

# **Painful intelligence: What AI can tell us about human suffering**

Aapo Hyvärinen

University of Helsinki

Version 1.0  
1st June 2022

## Abstract

This book uses the modern theory of artificial intelligence (AI) to understand human suffering or mental pain. Both humans and sophisticated AI agents process information about the world in order to achieve goals and obtain rewards, which is why AI can be used as a model of the human brain and mind. This book intends to make the theory accessible to a relatively general audience, requiring only some relevant scientific background.

The book starts with the assumption that suffering is mainly caused by frustration. Frustration means the failure of an agent (whether AI or human) to achieve a goal or a reward it wanted or expected. Frustration is inevitable because of the overwhelming complexity of the world, limited computational resources, and scarcity of good data. In particular, such limitations imply that an agent acting in the real world must cope with uncontrollability, unpredictability, and uncertainty, which all lead to frustration.

Fundamental in such modelling is the idea of learning, or adaptation to the environment. While AI uses machine learning, humans and animals adapt by a combination of evolutionary mechanisms and ordinary learning. Even frustration is fundamentally an error signal that the system uses for learning. This book explores various aspects and limitations of learning algorithms and their implications regarding suffering.

At the end of the book, the computational theory is used to derive various interventions or training methods that will reduce suffering in humans. The amount of frustration is expressed by a simple equation which indicates how it can be reduced. The ensuing interventions are very similar to those proposed by Buddhist and Stoic philosophy, and include mindfulness meditation. Therefore, this book can be interpreted as an exposition of a computational theory justifying why such philosophies and meditation reduce human suffering.

Copyright ©2022 Aapo Hyvärinen. All rights reserved.

Distribution allowed as per Creative Commons Attribution-Noncommercial-NoDerivatives (CC BY-NC-ND) License.

This book is typeset in Utopia(R), Copyright 1989, 1991 Adobe Systems Incorporated. All rights reserved.

Utopia is either a registered trademark or trademark of Adobe Systems Incorporated in the United States and/or other countries.

Used under license, as implemented in the LaTeX `fourier` package.

# Contents

<b>1</b>	<b>Introduction</b>	<b>8</b>
	Investigating intelligence by constructing it . . . . .	8
	Is the brain a big computer? . . . . .	9
	Machine learning as analogue to evolution . . . . .	10
	Can an AI actually suffer? . . . . .	10
	Intelligence is painful—overview of this book . . . . .	11
<b>I</b>	<b>Suffering as error signalling</b>	<b>14</b>
<b>2</b>	<b>Defining suffering</b>	<b>15</b>
	Medical definitions of pain . . . . .	15
	Medical and psychological definitions suffering . . . . .	16
	Ancient philosophical approaches to suffering . . . . .	19
	Two main kinds of suffering . . . . .	20
	Using the pain system for broadcasting errors . . . . .	20
<b>3</b>	<b>Frustration due to failed plan</b>	<b>22</b>
	Agents, states, and goals . . . . .	22
	Planning action sequences, and its great difficulty . . . . .	24
	Frustration as not reaching planned goal . . . . .	26
	Defining desire as a goal-suggesting mechanism . . . . .	27
	Intention as commitment to a goal . . . . .	28
	Heuristics can help in planning . . . . .	29
<b>4</b>	<b>Machine learning as minimization of errors</b>	<b>32</b>
	Neurons and neural networks . . . . .	32
	Finding the right function by learning . . . . .	35
	Learning as minimization of errors . . . . .	38
	Gradient optimization vs. evolution . . . . .	40
	Learning associations by Hebbian rule . . . . .	41
	Logic and symbols as an alternative approach . . . . .	42
	Emergence of unexpected behaviour . . . . .	45

<b>5 Frustration due to reward prediction error</b>	<b>47</b>
Maximizing rewards instead of reaching goals . . . . .	47
Learning to plan using state-values and action-values . . . . .	49
Frustration as reward loss and prediction error . . . . .	52
Expectations or predictions are crucial for frustration . . . . .	55
Unexpected implications of state-value computation . . . . .	57
Evolutionary rewards as obsessions . . . . .	58
Reward maximization is insatiable . . . . .	59
<b>6 Suffering due to self-needs</b>	<b>61</b>
Self as long-term performance evaluation . . . . .	61
Self as self-preservation and survival . . . . .	63
Self as desires based on internal rewards . . . . .	65
Uncertainty, unpredictability, and uncontrollability as internal frustration . . . . .	68
Fear, threat, and frustration . . . . .	69
<b>7 Fast and slow intelligence and their problems</b>	<b>72</b>
Fast and automated vs. slow and deliberative . . . . .	72
Neural network learning is slow, data-hungry, and inflexible . . . . .	74
Using planning and habits together . . . . .	76
Advantages of categories and symbols . . . . .	79
Categorization is fuzzy, uncertain, and arbitrary . . . . .	81
The many faces of frustration: Summarizing the mechanisms of suffering . . . . .	84
<b>II Origins of suffering: uncontrollability and uncertainty</b>	<b>88</b>
<b>8 Emotions and desires as interrupts</b>	<b>89</b>
Computation is one aspect of emotions . . . . .	90
Emotions interrupt ongoing processing . . . . .	91
Desire as an emotion and interrupt . . . . .	92
Emotions include hard-wired action sequences . . . . .	94
How interrupts increase suffering . . . . .	95
Emotions are boundedly rational . . . . .	98
<b>9 Thoughts wandering by default</b>	<b>99</b>
Wandering thoughts and the default-mode network . . . . .	99
Wandering thoughts as replay and planning . . . . .	101
Experience replay focuses on reinforcing events . . . . .	104
Replay exists in rats, humans, and machines . . . . .	106
Wandering thoughts multiply suffering . . . . .	107

<b>10 Perception as construction of the world</b>	<b>110</b>
Vision only seems to be effortless and certain . . . . .	110
Perception as unconscious inference . . . . .	113
Prior information can be learned . . . . .	115
Illusions as inference that goes wrong . . . . .	118
Attention as input selection . . . . .	119
Subjectivity and context-dependence of perception . . . . .	120
Reward loss as mere percept . . . . .	121
Ancient philosophers on perception . . . . .	122
<b>11 Distributed processing and no-self philosophy</b>	<b>124</b>
Are you really in control? . . . . .	124
Necessity of parallel and distributed processing . . . . .	126
Central executive and society of mind . . . . .	129
Control as mere percept of functionality . . . . .	132
Philosophy of no-self and no-doer . . . . .	134
<b>12 Consciousness as the ultimate illusion</b>	<b>136</b>
Information processing vs. subjective experience . . . . .	136
The computational function of human consciousness . . . . .	137
The origin of conscious experience . . . . .	140
Why is simulated suffering conscious? . . . . .	142
Self vs. consciousness . . . . .	144
Nothing is real? . . . . .	145
<b>III Liberation from suffering</b>	<b>148</b>
<b>13 Overview of the causes and mechanisms</b>	<b>149</b>
Why there is (so much) suffering . . . . .	149
Cognitive dynamics leading to suffering . . . . .	154
An equation to compute frustration . . . . .	157
<b>14 Reprogramming the brain to reduce suffering</b>	<b>160</b>
Reducing expectation of rewards . . . . .	161
Reducing certainty attributed to perception and concepts . . . . .	166
Reducing self-needs . . . . .	169
Reducing desire and aversion . . . . .	171
How far should reducing desires and expectations go? . . . . .	173
<b>15 Retraining neural networks by meditation</b>	<b>176</b>
Contemplation as active replay . . . . .	177
Mindfulness meditation as training from a new data set . . . . .	178
Speeding up the training . . . . .	181

Reducing interrupting desires . . . . .	184
Emptying the mind and reducing simulation . . . . .	186
Attitude of acceptance . . . . .	188
Metacognition and observing the nature of mind . . . . .	189
Letting go and relaxation as unifying principles . . . . .	192
<b>16 Epilogue</b>	<b>195</b>
<b>Index</b>	<b>200</b>
<b>Bibliography</b>	<b>207</b>

# Preface

I like to write books that I would have wanted to read myself as a student. I really wish I had been able to read this book. It would probably have changed my life and my career, as I would have insisted on doing my PhD on this topic. Alas, when I was a student in the 1990s, the topic of this book was not something a reasonable PhD student would have worked on. There was hardly any literature on the topic; it would have been considered uncharted territory, if not suspicious. I hope the world has changed, and that this book may contribute to that change. With the huge increase in research on AI and computational neuroscience on the one hand, and affective neuroscience and mindfulness meditation on the other, I think the time is ripe to attempt a synthesis, which is the motivation for this book.

What I should emphasize is that this book is about a scientific theory, or rather, several scientific theories. It is not a book that teaches meditation; it has little to do with self-help and certainly constitutes no clinical guidance. Nor is it really a philosophical book in the sense that the word would be used in academic circles: while there is some philosophical speculation, the main paradigm is that of the natural sciences. It may be surprising that I seem to include artificial intelligence in the natural sciences, but here it is largely used as a computational model of the brain, even if sometimes on a very abstract level. The strong neuroscience component of this book further connects it to empirical science.

I have tried to write the book so that it is suitable for as wide an audience as possible. I believe anybody trained in computer science or neuroscience should be able to understand it. Scientific training in any discipline might be enough to understand the main ideas, and I hope that even members of the general public might find something interesting in it. Although not primarily intended as such, the book can also be used as a university-level textbook for advanced undergraduates or graduate students in computer science or cognitive science; it should also be suitable for computationally minded students in neuroscience or psychology.

This book was written while working in different institutions. Most of the work was done while a faculty member at the University of Helsinki (Department of Computer Science). Part of the writing was accomplished while a faculty member at University College London (Gatsby Computational Neuroscience Unit) as well as a research scientist at Université Paris-Saclay (DataIA Institute and Inria-Saclay-Ile-de-France, supported by grant ANR-17-CONV-0003). The work was further supported by a Fellowship from CIFAR (Learning in Machines & Brains Program).

Finally, I'm very grateful to Moritz Grosse-Wentrup, Riitta Hari, Marianne Maertens, John Millar, Tiina Parviainen, Jonne Viljanen, and, especially, Michael Gutmann, for most helpful comments on the manuscript.

*Helsinki, June 2022*

Aapo Hyvärinen

# Chapter 1

## Introduction:

## Understanding human suffering by AI

What is the most central question in human life? For me, it is the question of suffering. There may be questions which are more fundamental, or philosophically more fascinating, for example: Why does the world exist? Or, how is it possible that we are conscious? But those questions are rather theoretical and mainly satisfy one's intellectual curiosity. If you found the answer to those latter questions, would that change your life, or other people's lives, for the better?

The question of suffering is with us at every moment. By suffering I mean mental pain, the opposite of pleasure and happiness. In some cases, it is a result of physical pain, but usually of purely mental origin. In fact, any casual observer of human life easily comes to the conclusion that it is full of such suffering: There is frustration, anxiety, sadness, depression, and so on.

Why is the "human condition" so unpleasant: did somebody (or something) make a huge mistake in designing humans? And, most importantly, is there anything we can do about it: can we remove suffering, or at least reduce it? Now, this is a question that has enormous practical significance. Reducing suffering, almost by definition, makes people's lives better.

The starting point of this book is the idea that we can use the theory of artificial intelligence, or AI, to understand why there is so much suffering in humans. This book will show how suffering is largely due to the inability of an intelligent system, whether an artificial intelligence or a human being, to understand its own programming and its own limitations, in particular the limitations of its computation and data.

### Investigating intelligence by constructing it

How can I claim that the theory of AI has any relevance to understanding the human mind, let alone suffering? The answer lies in how AI can help us understand the computational design principles which are applicable to humans as well.

When I asked above if somebody made a huge mistake in designing humans, that "somebody" was of course evolution, metaphorically speaking. Evolution designed the basic processes of our mental life, for good or bad. Importantly, evolution didn't construct our brains in some random, arbitrary ways, but it designed us to be fit for certain purposes and goals. Ultimately, those evolutionary goals are about reproduction and



spreading your genes, but to satisfy that ultimate goal, many more intermediate goals need to be considered. You have to get food, find sex partners, and not get killed. These, in turn, require that you know how to walk, and you are able to recognize objects as well as to plan your future actions.

We can learn to understand such evolutionary design goals by trying to design and construct an AI, or a robot. This is a perspective which is gaining more and more prominence in neuroscience: Trying to actually construct an intelligent system forces you to think about the computation and algorithms needed.

Ordinary neuroscience is based on conducting experiments on humans or animals. It can establish many interesting facts about the brain; for example, where in the brain the processing necessary for vision or fear takes place. It can also tell us something about how such processing happens; for example, by explaining how the brain recognizes that the animal in front of you is a cat and not a dog, and why we would get scared if it were a tiger.

However, the deepest question in neuroscience is the “why” question: Why does a certain kind of processing take place at all? What is its evolutionary purpose? Why do we, for example, have emotions like fear in the first place? Why is our mind frequently assailed by thoughts about the past and the future even when we try to concentrate on the present? And ultimately, why is there suffering?

Designing intelligent systems goes a long way toward answering the “why” question. If we find that an AI necessarily needs a certain kind of computation to achieve human-like intelligence, it is likely that the human brain does that same kind of computation—at least on some level of abstraction. AI can also give us a deeper understanding of “how” computations happen in the human brain, since designing it necessarily forces the scientists to figure out all the details needed in the computation.

## Is the brain a big computer?

The prerequisite for learning about the brain by building intelligent systems is that our brain is in many ways like a computer. In fact, the modern paradigm in neuroscience and psychology considers the brain as an information-processing device. The term “cognition” is used to describe information-processing performed by the brain, while with ordinary computers we usually talk about computation.

The brain receives new data by seeing, hearing, or otherwise sensing things. It processes the sensory data in various ways, ultimately enabling us to recognize objects and act in the world. It can also process information retrieved from its own memory, which is necessary for what we call thinking in plain English. A system that processes information in such ways can be called, almost by definition, a computer, so it is natural to say that, actually, the brain *is* a computer.

Certainly, the brain is very different from any ordinary computer that you can buy in a shop. For example, your PC, or your mobile phone, has a central processing unit (CPU), sometimes a couple of them. The brain has no such thing. The information-processing happens in the neural cells, or neurons. Each of them is like a tiny CPU which can only perform extremely simple processing— but there is a huge number of them, tens of billions. The crucial difference with respect to a CPU is that each neuron processes its own input independently, and all the neurons do that at the same time—this is called parallel and distributed processing.

Yet, from an abstract viewpoint, such differences can be seen as just technical details. In particular, if we are interested in the question of “why” certain computations are performed, the physical structure of the information-processing device, or even the details of the programming do not matter. What really matters for our purposes is whether the brain and the computer need to solve the same kinds of computational problems.

This will be the case if humans and the AI live in the same kind of environment, have the same kind of goals for their actions, and use similar means to try to reach them. That is increasingly the case when AI develops in terms of autonomous robots, for example, and in any case, we can use our current AI theory to extrapolate what AI's might be like in the future.

## Machine learning as analogue to evolution

Even granted that humans and computers are both information-processing devices, some would argue that they process information based on very different principles. A popular claim is that a computer does exactly what it is programmed to do, and nothing else, and this is supposed to be very different from humans who do what they want themselves—so any parallels between humans and computers are impossible. I think this reasoning is fundamentally wrong, for two reasons.

First, modern AI systems do not just do what they are programmed to do. That's because their function is based on *learning*. They are programmed to learn from input data. The input may be a database determined by the programmer; it can be obtained by crawling the internet; or it can be the result of interactions with the environment, like a robot using a camera or users typing words, and so on. What the programmer really does is to provide an *algorithm* for learning. The algorithm is based on certain goals or *objective functions* that the AI is trying to optimize. An AI dedicated to searching the internet for images that resemble a given target image will learn to optimize the accuracy of its search results, for example by maximizing the number of clicks users make on each image it proposes.

What this means is that anyone who programs an AI cannot really know in detail what the AI will actually do, because it is often impossible to know what kind of input the AI will receive, and it is equally difficult to understand what the AI will learn from it. Even in the simplest case where the programmer completely decides the input to the AI, the input is often so complex (say, millions of pictures downloaded from the internet) that it is impossible for a human programmer to understand what can be learned from that data.

The second reason why this is not a major difference between humans and AI is that just like an AI is programmed by humans, we humans are designed—one might say “programmed”—by evolution. From an evolutionary perspective, we are programmed to maximize an objective function which roughly is given by the total number of copies of our genes in the population. To satisfy such programming, we gather a lot of data—by reading things, talking to people, and simply looking around—which is not so different from an AI.

So, I have turned the claim about the difference between AI and humans on its head. What humans and AI have *in common* is that both are programmed *by something else* to have certain goals and needs; nobody has really decided “by themselves” to have the needs and goals they have. To accomplish those goals, both humans and AI gather data from the environment and learn from it, which leads to actions that are *very difficult to predict*. So, in the end there is little difference between AI and humans, except regarding the source of the original programming—whether it was by evolution or a human programmer.

## Can an AI actually suffer?

By now, I hope to have convinced you that an AI is a useful model of many phenomena taking place in the human brain. But perhaps there are limits. Some would argue that we cannot talk about AI's or robots suffering:

They may *seem* to be suffering, or look like they are suffering, but in fact they are not, because they cannot *feel* anything.

I think this argument may not be completely wrong, but it is quite irrelevant. Obviously, it depends on the exact definition of what suffering is. It is true that AI may not feel suffering in the same way as humans because that might require that it is conscious. This argument against AI's suffering really hinges on two points: First, that an AI is not conscious, and second, that consciousness is necessary for suffering.

However, conscious feeling is only one part of suffering. The situation is similar with emotions, such as fear, which are actually clever information-processing mechanisms. The conscious feeling of being afraid is only one part of a complicated process involving cognition (or information-processing), behavioural tendencies, and several other aspects. I would argue it is the same for suffering.

Suffering is actually a signal in a complex information-processing system. The real meaning of the suffering signal is that an *error* occurred—this will be elaborated in several chapters in this book. Any information-processing system can create error signals. That's why we can, in that specific sense, say that an AI or a robot is suffering, even if they are not conscious. All that would be missing is the conscious feeling components of suffering.

There is an even more important reason why it is largely irrelevant here if an AI really suffers according to some stringent definition of the word. This book does not just aim to describe the mechanisms of suffering; the primary goal here is to develop various ways of alleviating suffering. For the purpose of *reducing* suffering, it does not matter if computers actually suffer in some deep sense. If we can reduce suffering in an AI that is sufficiently human-like, then, with reasonable probability, the same methods will apply to humans, and they will reduce suffering in humans, including the conscious experience of suffering. In other words, the AI is really a simulation or a model of mechanisms that are relevant for making humans happier.

For those who find it impossible to think that a computer could suffer in any sense of the word, I suggest the following viewpoint that they can use while reading this book. Trying to understand human suffering by AI is one big thought experiment, where we try to understand how much the AI would suffer under various circumstances, *if* it were able to consciously experience suffering. It is like a mathematical model of atoms, or like a computer simulation of chemical processes. Everybody agrees that models and computer simulations are not the real thing, but they can help us understand the actual natural processes, and in particular, predict their behaviour. A model may tell you how a change in one quantity, say X, leads to a change in another quantity, Y. If you know that, you can perhaps choose X to maximize or minimize Y—which might be suffering.

## **Intelligence is painful—overview of this book**

The central claim in this book is that if we create an artificial intelligence that is really intelligent, really worthy of its name, it is necessarily going to suffer—more or less like humans. In spite of the many differences between AI's and humans, there is a common logic in the design. In order to achieve sufficiently human-like intelligence, certain design principles have to be followed, and these lead to suffering. This book explores several interwoven ideas about such a computational basis of suffering, and the necessity of suffering as a part of intelligence.

The fundamental principle is that suffering is caused by error signalling, which is typically due to frustration. Frustration occurs when the system, generally called an “agent”, fails to achieve a goal. Such errors are inevitable in a complex world, where things are uncertain and unpredictable, and we have limited control over them. Error signalling is necessary for any sufficiently intelligent system, since such error signals are used by

learning algorithms. Our brain produces error signals automatically, and we simply cannot shut off the error-signalling system.

In fact, the complexity of the world is overwhelming for any known intelligent system, whether the very best supercomputer in the world, or the most intelligent human brain. The computations available to them cannot handle all the different possibilities, for example, in choosing action sequences to reach a given goal.

Modern AI uses *learning* to cope with such complexity. Learning from complex data enables particularly sophisticated information processing. However, for such learning to be really successful, huge data sets are required. Obtaining data sets which completely capture the complexity of the world is rarely possible in practice.

These two factors, lack of computational resources together with scarcity of data, mean that the intelligent agent cannot work optimally. Its intelligence, and its control over the world, are limited. Thus, there will be errors: Its actions do not always lead to the desired outcome. This is the fundamental reason why such errors, and suffering, are ubiquitous.

Suffering is greatly enhanced by several information-processing principles inherent in the design of human-like intelligent systems. One is the phenomenon of experience replay, where memories related to past errors are recalled and repeated in the system in order to optimize learning about past experiences. Likewise, plans for future actions are constantly computed, which means the agent simulates or “imagines” them in its mind, together with the ensuing errors. Such replay and planning multiply any suffering arising from real events. Errors are signalled as if those bad, imagined events happened for real. Replay and planning even lead to conscious suffering in humans. Somehow, some part of the brain does not understand that recalling or imagining an event is not the same as actually living it. That is why we suffer from mishaps which only happen in our imagination.

Modern AI has found systems based on parallel and distributed information processing to be useful for programming intelligent systems, which makes it understandable that our brain uses similar principles. However, parallel and distributed systems are difficult to control by any central “executive”. Instead, different computational modules are competing for control, making, for example, any sustained attention or concentration difficult. Any internal control of the agent’s computations is further reduced by emotions such as fear, which work as evolutionarily conditioned “interrupts” of ongoing processing. Thus, the agent has little control of even its own internal processing, let alone the external world.

Yet another problem is the difficulty of understanding how uncertain most perceptions and inferences are. Perceptions are often highly subjective and contextual interpretations, sometimes little more than guesses. Yet, humans often mistakenly think that our perceptual systems are able to discover some underlying objective reality. To simplify the overwhelming complexity of the world, an agent may further divide it into categories. However, the categories may be arbitrary and assigning objects to categories difficult. Our inability to appreciate such uncertainty and even arbitrariness leads to more suffering.

Finally, the goals and desires that have been programmed in us by evolution are ultimately counterproductive and make us unhappy. Evolution never had our happiness as its goal anyway. In fact, it forces us to do things which are clearly bad for our happiness, something I call “evolutionary obsessions”. Evolution makes us worry about our survival and our evolutionary performance, creating a sense of self. In fact, evolution does *not* want us to reduce suffering because the error-signalling system is necessary for learning and optimal behaviour. Evolution *does* want us to learn to act in more and more efficient ways, but the goals towards which this intelligence is used are those set by evolution, not us. Even worse, both AI and humans are usually trying to satisfy their drives and desires endlessly, without any limits; at no point do they become satiated and think

that they have achieved enough.

However, there is hope. At the very end of the book, I sketch methods that can be used to decrease suffering, based on the theories outlined in this book. What is needed is a reprogramming of the brain. The key method is to retrain the brain by inputting new data into the learning system. The new data will change the computations in such a way that error signals, and in particular frustration, are reduced. This is difficult and takes a lot of time, but various forms of philosophical contemplation and meditation attempt to do it. These methods are rather logical consequences of the theory, while at the same time, they have mostly been proposed earlier in Buddhist, and to some extent Stoic, philosophy. Thus this book can be seen as an attempt to construct a scientific, computational theory on the underpinnings of such philosophies.

## **Part I**

# **Suffering as error signalling**

**The first part will explore the very definition of suffering,  
existing proposals on how suffering comes about,  
and how these can be understood by the theories of AI and evolution**

## Chapter 2

# Defining suffering

In this chapter, I try to define the word “suffering”. This is not an easy task, as we will quickly see. Defining the term properly requires, to some extent, elucidating the underlying mechanisms creating suffering.

One fundamental point here is that I exclude physical pain from the definition of suffering; I use the word suffering synonymously with mental pain. Nevertheless, I will start the search for a definition of suffering by considering the closely related concept of pain, taken here in the medical sense of physical pain.

The central conclusion of this chapter is that the main definitions of suffering consider it based on either *frustration* or a *threat to the intactness* of the person. These two definitions, and especially the definition based on frustration, are the basis of the developments of the rest of this book. From a more abstract viewpoint, I will argue that such suffering can be seen as *error signalling*, similarly to physical pain.

### Medical definitions of pain

Let us start by defining pain. Pain has been given a widely accepted consensus definition by the International Association for the Study of Pain (IASP) as

Pain is an unpleasant sensory and emotional experience associated with actual or potential tissue damage or described in terms of such damage.

Surprisingly, while this definition was originally adopted in 1979, it is still used with minimal modifications. It posits damage to any tissue of the person, or any threat of such damage, as the origin of pain. Pain is then defined as an ensuing unpleasant experience. While this definition has been found to be quite useful in a clinical context, deeper theoretical analyses have found various problems.<sup>1</sup>

One important controversy is whether one should define pain as a subjective experience, or as something that has a more objective existence. The definition above talks about an “experience” which is here interpreted as a conscious, subjective experience: something that only I am aware of, and which you cannot measure in any objective way. As we will discuss in more detail in Chapters 8 and 12, this problem of subjective conscious experience vs. objectively observed phenomena is ubiquitous in neuroscience and psychology.

---

<sup>1</sup>Cohen et al. (2018) gives a long review of competing definitions; Corns (2016) considers the validity of the very concept; Klein (2007) proposes an alternative definition and reviews some philosophical approaches. IASP has very recently proposed a revised version (Raja et al., 2020), but the changes are minimal and rather immaterial for our purposes.

The problem with talking about such subjective experience in a scientific context is that objective, reproducible measurement is the basis of science. Fortunately, subjective experience can be measured in various indirect ways, such as verbal report. That is, we can ask the patient if there is pain. Yet we will never know for sure what the patient actually feels. In particular, we cannot tell how her experience of pain compares with other people's experience: Does she feel more pain than some other patient?

This problem in the definition is to some extent alleviated by the reference to tissue damage, which is objectively measurable and reasonably well-defined. Yet, as this definition clearly points out, actual tissue damage is not *necessary* for pain—since it can be just “potential”—and thus it does not provide a basis for measuring pain or for objectively defining it. (In fact, the definition does not actually say that pain is in any sense proportional to the amount of damage— it is well-known that tissue damage that creates a lot of pain in one person may create little pain in another—which complicates any measurement even more.) Another related problem with the IASP definition above is that it relies heavily on the word “unpleasant”, which is not a very well-defined term, and, again, quite subjective.

One approach to solve these problems is to take an evolutionary approach. To begin with, we could replace “unpleasant experience” in the definition by “experience that has evolved to motivate behaviour, which avoids or minimises tissue damage, or promotes recovery”.<sup>2</sup> Here we go towards defining pain using its evolutionary, functional role, while still acknowledging the subjective nature of pain by talking about an “experience”. The downside of such an approach is that it works on a very abstract level, and provides no details on what might cause pain, in contrast to the IASP definition which explicitly points at tissue damage (even if only potential). This definition, in a sense, shifts the burden to understanding the evolutionary goals of certain experiences, which is not easy either. However, one obvious candidate for such an evolutionary goal is minimizing tissue damage and recovering from it, which links this evolutionary approach with the IASP definition. In more general terms, the evolutionary goal could be the maintenance of “homeostasis”, that is, an optimal balance in the physiological condition of the body.<sup>3</sup> Such evolutionary logic can be applied on suffering as well, and we will see related argumentation throughout this book.

## Medical and psychological definitions suffering

In contrast to pain, suffering is a rather neglected term in science, and there is nothing like a consensus definition. Intuitively, most people would think suffering also contains an *unpleasant feeling or experience* as an integral part, while being more abstract and general than physical pain, in particular including more psychological and emotional aspects. A typical dictionary definition is “Feeling of pain or strong stress, either physical or emotional”.<sup>4</sup> Like pain, suffering is often considered a subjective experience which cannot be objectively measured.<sup>5</sup>

One simple and concrete approach to define suffering is to give examples of phenomena related to suffering and possibly producing suffering. A typical list would contain grief, sadness, discomfort, distress, anguish, fear—which is just a random sample, and many different lists can be produced. While this is a good starting

---

<sup>2</sup>This definition is by Wright (2011). On a related note, Seymour (2019) emphasizes the importance of pain as a signal used in control and learning, and relativizes the importance of conscious experience.

<sup>3</sup>(Craig, 2003)

<sup>4</sup><https://psychologydictionary.org/suffering/>

<sup>5</sup>(Cassell, 1982; Edwards, 2003)



point, it does not lead to a solid scientific theory.

Terms such as psychological pain or mental pain are often preferred in neuroscience, and some attempts at definitions of those terms have been made.<sup>6</sup> In this line of thinking, suffering is really a generalization of pain. This may not solve the problem of defining suffering, since the burden is then simply shifted to defining pain, but then we can leverage the large literature on pain, in particular the IASP definition just given, as well as any of its critique and improvements.

One approach distinguishes three kinds of pain: physical pain, social pain, and psychological pain.<sup>7</sup> An interesting emphasis in this line of research is that all these different kinds of pain are neurally very similar in the sense that the brain areas responsible are the same.<sup>8</sup> Here, physical pain is primarily due to physical damage to the body, but it can also be felt when there is a strong anticipation of such physical damage (think about going to a dentist), reminiscent of the IASP definition. In contrast, social pain is an unpleasant feeling due to social exclusion or rejection. Psychological pain is largely the same as what I call mental pain or suffering, and attempt to define here.

Importantly for our purposes, in such an approach, mental or psychological pain is often assumed to be due to *reward loss*, defined as follows<sup>9</sup>

[Reward loss is] a negative discrepancy between expected and obtained rewards.

In other words, reward loss happens when you expect a reward but don't get it, and it leads to mental pain. Reward loss can also be called *frustration*, although sometimes this term is reserved for the actual suffering caused by reward loss. This provides one important computational viewpoint: reward loss is a function of computations involving expectations, observations of the obtained reward, and their difference.

An alternative approach emphasizes how suffering is related to our person, or self. Psychological or mental pain has been characterized as an aversive state of high self-awareness of inadequacy,<sup>10</sup> or a negative appraisal of an inability or deficiency of the self.<sup>11</sup> This is analogue to IASP definition of physical pain in the sense that there is "damage", but on a more mental level, to our image of ourselves as a psychological and social entity.<sup>12</sup>

A particularly potent and influential idea in this vein is that suffering necessarily involves a threat to, or a loss of, the *intactness of the person*, as proposed by Cassell:<sup>13</sup>

Suffering is a state of severe distress induced by the loss of the intactness of person, or by a threat that the person believes will result in the loss of his or her intactness.

<sup>6</sup>Reviews on the topic are provided by Mee et al. (2006); Tossani (2013); Papini et al. (2015). The term "mental pain" could be criticized because all pain is ultimately mental, as seen in the IASP definition. In this book, I mainly use the term "suffering".

<sup>7</sup>(Papini et al., 2015; Eisenberger and Lieberman, 2004; MacDonald, 2009). Pain based on empathy when one sees others hurting, or "vicarious" pain, could be added to the list (Singer et al., 2004).

<sup>8</sup>However, see Iannetti et al. (2013); Wager et al. (2016) for criticism of the reverse inference used in that work. Iannetti and Mouraux (2010) argue that the brain network considered may be more related to detection of saliency (i.e. how much attention a stimulus attracts).

<sup>9</sup>(Papini et al., 2015)

<sup>10</sup>(Baumeister, 1990; Orbach et al., 2003)

<sup>11</sup>(Meerwijk and Weiss, 2011)

<sup>12</sup>In this line of research, typical in the philosophy of medicine and bioethics, suffering is sometimes seen as something particularly strong (Degrazia, 1998; Hoffmaster, 2014), in particular stronger than any pain typically encountered in everyday life. I don't follow such a definition here: Suffering can be very mild or very strong.

<sup>13</sup>(Cassell, 1989)

This is a natural generalization and abstraction of the IASP definition of pain as related to “tissue damage”. Damage to the intactness of the person includes tissue damage but is something much more general, in particular including damage to one’s self-image. It is of course crucial here to understand what “intactness” means more precisely; Cassell emphasizes the generality of this notion, saying that “suffering may occur in relation to any aspect of personhood”.<sup>14</sup>

A unifying theory, which combines several different kinds of suffering and pain in a single framework by linking them through the concept of frustration, has been developed by van Hooft.<sup>15</sup> He starts from an Aristotelian conception of the human person as having four “parts of the soul”. They range from the lowest level of biological functioning to the emotional/desiring functions and the rational functions, finally reaching the sense of the meaning of existence. In his theory, each of these parts has its own goals, its own form of “fulfillment”, which is again an Aristotelian idea. Suffering is then nothing else than frustration, namely “frustration of the tendency towards fulfillment” of one of the different parts of the soul. In this theory, the lowest level of biological functioning is even below ordinary pain and pleasure, and simply about staying healthy and alive. Ordinary physical pain is the frustration on the emotional/desiring level, where the goal of the organism is to gain pleasure and avoid pain. Frustration of rational (intellectual) function refers to suffering which happens when it is not possible to reach long-term goals that one would usually expect to reach and plan for. Frustration on the highest, “spiritual” level happens when it is impossible to understand why it is me that is sick—in the medical context where van Hooft writes—or life seems meaningless due to the despair and fear which a malady brings with it. This last kind of suffering brings us close to the kind of suffering considered in existential philosophy.<sup>16</sup>

Closely related definitions can be found in the literature on *stress*: Lazarus and collaborators define “psychological stress” as “a particular relationship between the person and the environment that is appraised by the person as taxing or exceeding his or her resources and endangering his or her well-being”.<sup>17</sup> So, we have to consider the possibility that stress is another kind of suffering, or a mechanism for suffering. However, I don’t take such a view in this book because the classic definition by Hans Selye, “the father of stress”, proposes that “stress is the non-specific response of the body to *any* demand” (my italics). This is a very general definition, and Selye has explicitly emphasized that positive, happy events can induce stress just as well as negative, threatening ones; think about an athlete engaged in a competition. Based on this definition, it does not seem possible to simply consider stress as one kind of suffering, unless we focus on the negative kind of stress, termed “distress” by Selye.<sup>18</sup> The distinction between distress and “pleasant” stress is, unfortunately, not very clear; it has been proposed that it is the *unpredictability and uncontrollability* of a situation which distinguish the unpleasant distress from other kinds of stress.<sup>19</sup> Their connection to suffering will be considered from different viewpoints in this book.

---

<sup>14</sup>For recent critique of Cassell’s approach, see Bueno-Gómez (2017) who criticizes Cassell’s definition precisely on the ground that “intactness” is not well-defined and may not exist; another point of critique is that Cassell’s definition ignores existential suffering. Further criticism is given by Tate and Pearlman (2019) who propose to define suffering as “a loss of a person’s sense of self” together with “a negative affective experience”.

<sup>15</sup>(Van Hooft, 1998)

<sup>16</sup>(Svenaeus, 2014; Bueno-Gómez, 2017)

<sup>17</sup>(Lazarus and Folkman, 1984); see also Lazarus (1993). Their work emphasizes the individual’s perception and interpretation of the events by the term “appraise”, related to Cassell’s definition which talks about “believing”. Another related approach to defining stress emphasizes conservation of resources, and defines the stress as, roughly, loss of resources (Hobfoll, 1989).

<sup>18</sup>See Fink (2017), where the quote from Selye is also taken from.

<sup>19</sup>(Koolhaas et al., 2011)

## Ancient philosophical approaches to suffering

Centuries before any such modern developments, some ancient philosophers already made great progress in understanding suffering. The best expert on the topic may have been the Buddha, and in fact the whole of Buddhist philosophy can be seen as a theory of suffering—especially when considering the original version proposed by the Buddha himself. He gave the following description of suffering:<sup>20</sup>

Union with what is displeasing is suffering; separation from what is pleasing is suffering; not to get what one wants is suffering.

This is actually not so much a definition of what suffering is, but rather an attempt to describe what the main causes of suffering are.

Stoic philosophers in ancient Greece and Rome had very similar ideas. Epictetus, one of the most famous Stoics, describes mechanisms that lead to suffering as follows:<sup>21</sup>

[D]esire promises the attainment of that of which you are desirous; and aversion promises the avoiding that to which you are averse. However, he who fails to obtain the object of his desire is disappointed, and he who incurs the object of his aversion wretched.

These are essentially a reformulation of the points given by the Buddha above. We can summarize these philosophical ideas as the following two causes for suffering, each with two variants:<sup>22</sup>

- 1 a) Not getting what one wants (Buddha, Epictetus)
- b) Something pleasant, which one would like to be present, is absent (Buddha)<sup>23</sup>
- 2 a) Not being able to avoid what one is averse to, i.e., wants to avoid (Epictetus)

---

<sup>20</sup>This is from a fundamental discourse by the Buddha found in one of the earliest known layers of Buddhist literature, the Pali Canon. Different versions are available in *Samyutta Nikaya* 56.11, *Majjhima Nikaya* 141, and *Digha Nikaya* 22, where the last one is the most detailed version. This quote is part of the description of what is called the Four Noble Truths, of which we here consider only the first one (see footnote 29 in Chapter 14 for the rest). The whole description of the first truth, synthesizing the different versions, says approximately: Birth is suffering, ageing is suffering, illness is suffering, death is suffering; grief, lamentation, pain, distress, and despair are suffering; union with what is displeasing is suffering, separation from what is pleasing is suffering, not to get what one wants is suffering. (Several partial translations of the Pali Canon are available on the internet and I will often select the translation I find the most compatible with my terminology; the one in the main text here is by Bhikkhu Bodhi.)

<sup>21</sup>Paragraph 2 in *The Enchiridion*, compiled approximately 125-135 CE. Quotes in this book are taken from the translation by E. Carter at [classics.mit.edu/Epictetus/epicench.html](http://classics.mit.edu/Epictetus/epicench.html).

<sup>22</sup>My logic is that if something pleasant is not present as in 1b, the point is that one actually *wants* that pleasant thing to be present, so this is also a question of not getting what one wants, as in 1a. The same logic shows that 2a and 2b are really the same thing. The points 1b and 2b present the difficulty that they use the terms pleasant (or “pleasing” in the translation quoted above) and unpleasant (or “displeasing”), much like the IASP definition of pain. (Alternative translations of these two words include “beloved”/“unbeloved” (Thanissaro Bhikkhu) “loved”/“loathed” (Nanamoli), “liked”/“disliked” (P. Harvey), and indeed “pleasant”/“unpleasant” (Piyadassi Thera), given at <https://www.accesstoinsight.org/tipitaka/sn/index.html#sn56>.) I suggest the key here is that “pleasant” is here assumed to necessarily lead to wanting (and “unpleasant” to aversion), and thus the Buddha is really talking about desire or wanting and aversion. When he specifically mentions wanting at the end of the quote, that may be seen as a kind of summary of the two first sentences.

<sup>23</sup>I interpret “separation”, also translated as “dissociation”, in the quote by the Buddha not simply as absence but as absence of something one would like to be there since it is pleasant.

- b) Something unpleasant is present (Buddha)

Then, the definitions by the Buddha and Epictetus can be interpreted in terms of wanting (point 1) and aversion (point 2) only. Point 1 in particular defines the typical case of frustration, related to reward loss considered above. Using the term somewhat liberally, we can also call the suffering in point 2 frustration, since the desire to avoid something is frustrated. Thus, we see that both the ideas of both the Buddha and Epictetus can be simply summarized as saying that suffering comes from frustration.

## Two main kinds of suffering

Now I shall try to recapitulate the ideas above, both ancient and modern, as succinctly as possible. I think we only need to talk about two kinds of suffering, or rather two mechanisms producing suffering, namely:

1. Frustration (e.g., Buddha, Epictetus, several neuroscientists<sup>24</sup>)
2. Threat to the intactness of the person, including their self-image (e.g., IASP, Cassell)

Furthermore, van Hooft's theory could be seen as combining these two aspects.

Based on this dichotomy, this book will develop two computational definitions of suffering, one each for these two aspects. Frustration will be specifically dealt with in Chapters 3 and 5, and the threat to the person (or "self") in Chapter 6. Chapter 6 will also propose how the two aspects can be seen in a unified framework by seeing threats to the person as frustration of certain long-term goals, a bit in the same sense as van Hooft's theory.

The emphasis in the following chapters is, obviously, on information processing. As already argued in the introduction, my main justification for talking about information-processing is that the framework of information-processing is a *practically useful* way of describing suffering in the precise sense that it can tell us something about how to *reduce* suffering. Information-processing is something that we can influence, something we can intervene on, so from a practical viewpoint, it is a very important aspect of suffering to investigate. Focusing on information-processing is also perfectly in line with the current emphasis on cognition in neuroscience and psychology; I see cognition as synonymous with information-processing.

## Using the pain system for broadcasting errors

To conclude this chapter, I discuss some computational principles that explain why pain and suffering are so closely related. First, I propose that on a more abstract computational level, both pain and suffering are essentially *error signals*, messages that something is going wrong from the viewpoint of the goals and rewards of the system. Clearly, frustration signals that something went wrong in terms of not getting what one wants, and a similar case will be made for the threat to the person in Chapter 6. Such error signals are in fact ubiquitous in artificial intelligence, where, in particular, they can be used for learning to choose actions better in view of maximizing rewards. We will see several kinds of error signals in the following chapters, and see how some of them can be interpreted in terms of suffering.

---

<sup>24</sup>Among recent work, see especially Papini et al. (2015), but the idea has a long history in experimental psychology as reviewed by Papini et al.

Pain is thus an evolutionarily primitive form of an error signal. Its unique feature is that pain signals are *broadcast widely* in the information-processing system. This is important in an agent whose computation is distributed into different modules (whether processors or brain regions). For such an agent, it is necessary that any really important signal uses a special pathway that allows it to be broadcast to all, or most of, the modules. The pain signal is indeed broadcast widely to different neural systems, and the signal can change the behaviour of the whole organism in terms of making it stop whatever it is doing and pay close attention to the pain. Furthermore, when an error signal drives the learning of the system, as we will consider in later chapters, it often needs to be observed by several of the modules, and such broadcasting is essential.<sup>25</sup>

Suffering is largely using the neural systems originally developed for physical pain, as already mentioned. This makes evolutionary sense if we think that computationally more sophisticated forms of error signalling, such as frustration, simply started using the evolutionarily older pain signalling pathway, adapting it for their own purposes. That was practical because the pain system already existed, and served well the purpose of broadcasting error signals to many brain regions. Using the physical pain system for signalling mental pain is thus a useful computational shortcut.<sup>26</sup>

Yet, merely talking about information-processing, as in a computer, may seem a rather incomplete description of suffering. Why does suffering *hurt*, if it is merely a signal in an information-processing system? This is in fact exactly the same problem that we encountered with the IASP definition of pain above: Is it a subjective experience, or something more objective and measurable?

The evolutionary rationale just described explains why suffering “hurts” in the same way as physical pain: the physical pain system is hijacked for the purposes of suffering or mental pain. (Perhaps this explains why we talk about mental “pain” in the first place.) The very dichotomy of experience vs. objective measurements is thus exactly the same for pain and suffering, since it is a question of similar experiences and neural pathways. Nevertheless, explaining why physical pain actually subjectively feels like it does in the first place, is an extremely difficult question; it is intimately related to the question of consciousness, which we defer to Chapter 12. We shall rather continue, in the next chapter, by elucidating the computational underpinnings of a particular form of error signal: frustration.

---

<sup>25</sup>The broadcasting hypothesis is closely related to the global workspace theory by Baars (1997), which will be treated in Chapter 12. However, while Baars links broadcasting it to consciousness, I think the broadcasting does not have to be conscious, especially in the case of pain or suffering. The hypothesis is also closely related to the earlier interrupt theory of emotions explained in Chapter 8. The broadcasting might happen through several specific connections between brain areas, or through a central hub.

<sup>26</sup>Such evolutionary arguments for using the same system were proposed by Eisenberger and Lieberman (2004), see also Papini et al. (2015). I am slightly confounding “pathways” and “systems” here: While the existing evidence is mainly about overlapping activation of certain brain regions, I am extrapolating the idea to the case of the signalling pathways.

## Chapter 3

# Frustration due to failed plan

In this chapter, I propose a first model where frustration is a fundamental mechanism for suffering. It is assumed that an agent, whether a human or an AI, engages in *planning* of action sequences in order to get to a desired *goal state*. A state is here an abstraction of the properties such as location, context, and possessions of the agent. Frustration happens when the goal state is not reached in spite of the agent executing the planned sequence of actions.

I start by emphasizing the great computational difficulty of such planning of action; it is one reason why frustration happens. Another central concept here is wanting or *desire*, which is a complex phenomenon we will return to several times in this book. As an initial definition, I consider desire as a computational process that suggests goals for the planning system. Finally, I discuss the importance of committing to a single plan, even in the presence of conflicting desires, based on Bratman's concept of *intention*. This chapter lays out the framework in simple, largely intuitive terms; the main terms and concepts will be greatly refined in later chapters.

### Agents, states, and goals

One may be tempted to think of an artificial intelligence as a system which just takes input, and processes information. However, information-processing in itself will actually be rather pointless unless it leads to some kind of visible output or action regarding the external world. In the very simplest case, action can just mean printing some text on a computer screen, so this is not necessarily a big leap.

In AI, the basic unit of analysis is often what is called an intelligent *agent*, i.e. a system which not only processes information but also takes actions. In fact, the word “agent” literally means “one that acts”. An intelligent agent can be artificial, such as a robot or an AI program, but the term also encompasses biological agents, that is, animals. In one extreme, an artificial agent could be just a program inside a computer, working in a virtual world with no physical body; actions would essentially consist of sending messages inside an information network. In the other extreme of human-like artificial agents, it could be a robot having a body with arms and legs; actions would include walking and grasping objects. In this book, we will see examples of both extremes—in addition to agents that actually are animals or humans.

Such an agent needs at least two things: perception and action selection. Perception is actually a tremendously difficult task but we defer its discussion to Chapter 4 and especially Chapter 10. To begin with, we assume perception is somehow satisfactorily performed, and consider the question of how the agent is to choose

its actions.<sup>1</sup>

Perhaps the simplest and most intuitive approach to action selection is to think in terms of *goals*. This is an introspectively compelling approach: We usually think of ourselves as acting because there are certain goals that we try to reach. That may be why it has also been a dominant approach in the history of AI, starting from around 1960. For example, a very simple thermostat has the goal of keeping room temperature constant; a cleaning robot has the goal of removing dust and dirt from the room.

The goals of humans are fundamentally determined by evolution, complemented by societal and cultural influences. In this book, I simply use the word “evolution” for brevity to describe the joint effect of biological evolution, culture and society. The assumption here is that the latter two are ultimately derived from biological evolution, although this is of course a controversial point. Fortunately, for the purposes of this book, the exact relationship between biology and culture is irrelevant: What matters is that the goals that humans strive for are largely, even if sometimes very indirectly, determined by some outside forces. Humans can set some intermediate goals, such as getting a job, but those are usually in the service of final biological or societal goals, such as being nourished or raising one’s social status. In the case of AI, in contrast, the goals are usually supplied by its human designers. This may seem to be a fundamental difference between AI agents and humans, but we will see in later chapters that it may not matter very much; the human designer plays a role similar to evolution in terms of being an outside force. In any case, regardless of where the goals come from, the way they are translated into action may still be rather similar in both cases.

### *Modelling the world as states*

In order to choose its actions, the agent should have some kind of a model of how the world works, where the agent itself is seen as a part of the “world” modelled. The model expresses the agent’s beliefs of what the world is typically like, and how the world changes from one moment to another, in particular as a function of the actions the agent takes.

AI research uses a very abstract kind of a world model based on the concept of a *state*, where each possible configuration of the world is one state. For example, if a cleaning robot is in the corner of a room, facing south, and there is only a single speck of dust in the room, at exactly two meters east from the robot, that is one state, we can call it state #1. If the agent finds itself 10 cm further to the west, it is in another state, say state #2; likewise, if another speck of dust appears in the room, that means the agent is in state #3 (and if the speck of dust appears *and* the agent is 10 cm further to the west, that is yet another state). In the simplest case, such a world model has states which are categorical, or discrete; in other words, there is a finite number of possible states. This is a very classical AI approach, but we will see alternatives in later chapters.

Any effects of the agent’s actions can now be described in terms of moving from one state to another, called state transitions. Indeed, in addition to knowing what the states of the world are like, the agent should know something about the transitions between the states caused by its actions. If it finds itself in state #1 and decides to move forward, does it find itself in state #2, or #47, or something else? If the agent’s world model can predict the effects of its actions in terms of transitions from one state to another, it is ready to start taking actions.

Using this formalism of states, the basic approach to action selection is that one of the states is designated as the goal state, by some mechanisms to be specified. The agent then uses its capabilities to reach the goal state, starting from its current state. That is surprisingly difficult, usually requiring complicated computation

---

<sup>1</sup>For standard textbooks on the topic, see (Russell and Norvig, 2020; Poole and Mackworth, 2010).

called *planning*, which is a central concept we consider next.<sup>2</sup>

## Planning action sequences, and its great difficulty

The fundamental problem in action selection is that you must actually select *sequences* of short actions. For example, if the agent in question is you, and you decide to get something from the fridge in the kitchen, you need to take a step with your left foot, take a step with your right foot, repeatedly, until you finally can open the fridge door—which consists of several actions such as: raise your arm, grab the handle, pull it down, pull the door, and so on. In some cases, you may easily know how to choose the right sequence, but it is not easy at all in many cases. For programming AI, this has turned out to be quite a challenge.

This is known as the problem of planning in AI. Using the formalism of world states, action sequences can be represented graphically as what is called a tree (Figure 3.1). The “root” of the tree represents the state you’re in at the moment. Any action leads to a branching of the tree, and depending on the action, you will find yourself on any of the new branches. (In this figure, taking an action means moving down in the tree, and we assume for simplicity that there are just two actions you can take at any time point). At the end of a given number of actions, or levels in the tree, you find yourselves at one of those states which are depicted at the outer “leaves” of the tree. Of course, the tree continues almost forever since you can take new actions all the time, but to keep things manageable, we consider a tree of a limited depth.

Let’s now assume that the agent has been given a goal state by the programmer. It would be one of the states at the lowest level of the tree. The central concept here is *tree search*; many classical AI theories see intelligence as a search for paths, or action sequences, among a huge number of possible paths in the action tree. In particular, the planning system finds a path which leads from the current state to the goal state. Such search may look simple, but the problem is that with such paths or action sequences, the number of possibilities grows exponentially. If you have, at any one time point, just two different actions to choose from, then after 30 such time points you have more than a billion (more exactly  $2$  to the power of  $30$ ) possible action sequences to choose from. What’s worse is that typically an AI would have many more than just two possible courses of action at any one point. The computations involved easily go beyond the capacity of even the biggest computers or brains. So, it may be impossible to “look ahead” more than a couple of steps in time.

The difficulty of such planning may be difficult for humans to understand since evolution has provided various tricks and algorithms that solve the problem quite well, as we will see below. We may only be able to grasp the difficulty of planning in some slightly artificial examples such as the search tree above. One of the more realistic examples would be planning a route between two points. Say you find yourself in a random location in Paris and want to go to the Eiffel Tower using public transportation. Even if you remembered every detail of the metro map as well as the geography of Paris itself, you would still need quite a lot of computation, in the form we would usually call “thinking”. Which metro station should I walk to, or should I perhaps use the

---

<sup>2</sup>It may seem very abstract to consider the world in terms of states. More insight might be obtained by taking an object-oriented viewpoint and considering the world a collection of objects. However, such a theory is still in its infancy (Diuk et al., 2008; Guestrin et al., 2003), so we have to use the approach based on world states. Another question is whether the states should really be considered discrete-valued, such as indexed by integers. In fact, most current AI systems do not use such a categorical representation, but rather some kind of continuous-valued perceptual representation in a neural network, for example given by the outputs of certain neurons. However, the approach using discrete states is widely used in theory and textbooks because of its conceptual simplicity. This is discussed in more detail in Chapters 4 and 7.



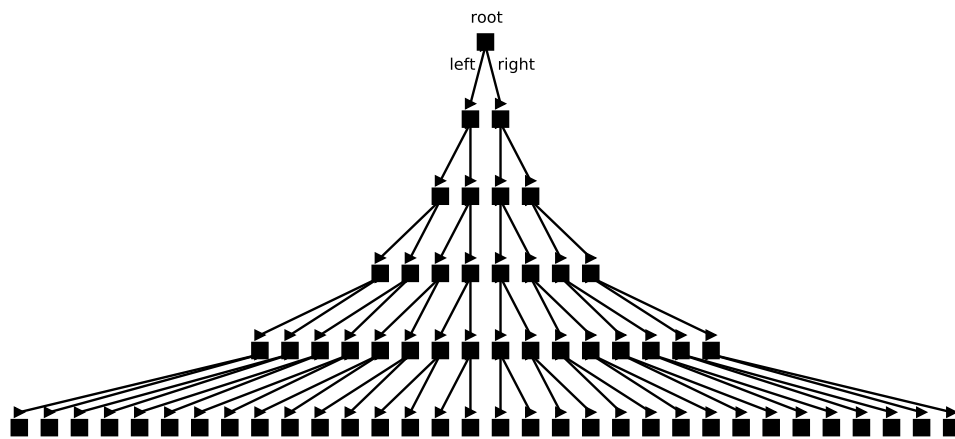


Figure 3.1: A search tree where the agent has two action options at every time point. They could be “turn left” or “turn right”, supposing the agent always makes a new decision when it finds itself in a new crossroads in a maze. The squares represent different states the agent can find itself in; the agent starts at the upper-most square in the figure (called root), and each action takes the agent one level down in this figure. The lines with arrows are the transitions to new states after every action taken. The crucial point here is that the number of different paths it can take grows exponentially. After just 5 steps, as depicted here, the number of paths equals 32 (2 to the 5th power). After 30 steps, it would be more than a billion.

bus? What is the best itinerary once inside the metro station? It is not surprising that people tend to use mobile phone apps to solve this problem.<sup>3</sup>

Board games are an extreme example of the difficulty of planning. Humans playing chess have great difficulties in thinking more than one or two moves ahead. The search tree has a lot of branches at every move because there are so many moves you can take. Even worse, your opponent can do many different things. (The uncertainty regarding what your opponent will do further adds to the complexity, but that is another story.)

A lot of the activity we would casually call *thinking* is actually some kind of planning. If you are thinking about where to go shopping for a new electronic gizmo, or how to reply to a difficult message from your friend, you are considering different courses of action. Basically, you're going through some of the paths in the search tree. Often, such thinking or planning happens quite involuntarily, even when you're supposed to be doing something else, a topic to which we will return in Chapter 9.

## Frustration as not reaching planned goal

Equipped with this basic framework for action selection, we are ready to define frustration in its most basic form. We start by considering one part of the Buddha's definition of suffering mentioned above: "not to get what one wants": This is in fact a typical dictionary definition of frustration. To achieve a deeper computational understanding of the phenomenon, we integrate this with the framework of planning.

Just like AI, complex organisms such as humans engage in planning: Based on their perception of the current environment, they try to achieve various goals by some kind of tree search. For such organisms, it is vital to know if a plan failed, so that they can re-plan their behaviour, and even learn to plan better in the future, as will be considered in detail in the following chapters.

We thus formulate the basic case of frustration as *not reaching a goal that one had planned for*, and the ensuing error signal. Sometimes frustration rather refers to the resulting unpleasant mental state; that is, frustration refers to the actual suffering instead of the cause for suffering. I use frustration in both meanings in this book.<sup>4</sup>

This initial definition will be refined and generalized in later chapters. In later chapters, we will see how central error signals are to any kind of learning. For example, a neural network that learns to classify inputs, or predict the future, is essentially minimizing an objective function which gives the error in such classification or prediction. Frustration can be seen as a special case of such error signalling: It signals that an action plan failed. In complex organisms like humans, which are constantly engaged in planning, frustration is an extremely important learning signal, and the basis of a large part of the suffering.

<sup>3</sup>Planning might actually seem to be very easy in a simple illustration like in Figure 3.1, since all you need is to *start* at the *goal* state, and go *backwards* in the search tree until you arrive at the root; thus you have found the path from the root to the goal. The reason why this does not work in practice is that in reality there are many overlapping trees, each starting from a *different* root state, and each goal state can be reached starting from a number of different roots. So, you cannot go backwards because you don't know which tree to follow. You can see this in the example of planning a route between two points in Paris: It may help a bit to start calculating backwards from Eiffel Tower, but you cannot just backtrack in a tree because the possible routes going "back" from the Eiffel Tower are as numerous as the routes you can start going "forward" from your current location; routes computed "backwards" from the Eiffel Tower can take you anywhere in Paris, not just your current location.

<sup>4</sup>This ambiguity is to some extent justified by the ambiguity of how the term is used in the literature, and some dictionaries explicitly list these two meanings for the term, e.g. <https://psychologydictionary.org/frustration/>.

## Defining desire as a goal-suggesting mechanism

However, there is a slight inconsistency here: Frustration was actually defined as not getting what one *wants* in Chapter 2. How is this related to our computational formulation based on planning above? In other words, what exactly is wanting, or desire, in a computational framework like ours?

In everyday intuitive thinking, action selection is indeed supposed to be based on wanting, or desires: An agent takes an action because it *wants* something, and it thinks it is reasonably likely to achieve or obtain it by that action. I choose to go to the fridge because I want orange juice. However, the account earlier in this chapter made no reference to the concepts of desire or wanting. In AI, the term “desire”, which I consider synonymous with “wanting”, can actually be used in a couple of different meanings.

In the very simplest definition, if the agent has a goal to plan for, one could simply say the agent “wants” to reach the goal state; desires would essentially be the same as goals. In such a meaning, desire is a kind of purely rational, “cold” evaluation of states and objects. However, the word has many more connotations in everyday language. Desire also has an affective aspect we could call “hot”, in which we are “burning with desire”, unable to resist it.

A definition that is a bit more in the direction of “hot” can be obtained by considering desire as a specific computational process inside the agent. To begin with, desire has been defined as a “psychological state of motivation for a specific stimulus or experience that is anticipated to be rewarding”.<sup>5</sup> While a “psychological state” may mean different things, here we consider a state as something where a particular kind of information processing is performed—another meaning would be related to conscious experience which we treat in Chapter 12. In practical terms, desire is often triggered by the perception of something that is rewarding to possess.<sup>6</sup> Such perception of an object often means that the agent should be able to get the object after a rather short and uncomplicated action sequence: If you see something, it is likely to be within reach.

From the viewpoint of information-processing, we thus define desire as: *A computational process suggesting as the goal a state that is anticipated to be rewarding and seems sufficiently easily attainable from the current state.* I want to emphasize that I am considering desire as a particular form of information-processing: Desire is not simply about preferring chocolate to beetroot, nor is it merely an abstract explanation of the behaviour where I grab a chocolate bar. It is sophisticated *computation* that is one step in the highly complex process that translates preferences into planning and, finally, into action.

The starting point for that processing is that your perceptual system, together with further computations, estimates that from the current state, you can relatively easily get into a state of high reward—the exact formalism for “rewards” will be introduced in Chapter 5. This realization will trigger, if you are properly programmed, further computational processes that will try to get you in that desired state by suggesting it as the goal for your planning system. When all this happens, you want the new state, or have a desire for that new state, according to the definition just given. For example, if chocolate appears in your visual field, your brain will compute that the state where you possess the chocolate is relatively easy to reach, and produces high reward; so it will choose the chocolate-possessing state as a possible goal and input that to the planning system.

Therefore, the definition of desire just given shows how the intuitive definition of frustration as not getting

<sup>5</sup>(Papies and Barsalou, 2015). For a number of alternative definitions see (Schroeder, 2017).

<sup>6</sup>“[D]esire arises when an internal or external cue triggers a simulation, or partial re-enactment, of an earlier appetitive experience that was rewarding.”(Papies and Barsalou, 2015). I’m using “reward” in this chapter in a non-technical sense, which will be refined in Chapter 5.

what one wants and the computational definition of not reaching the goal are essentially the same thing. This definition also solves a question which many readers must have asked while reading this chapter: Where do the goals for planning come from? In a very simple AI, there might be just a single goal, or a small set, defined by the programmer. But for a sophisticated agent, that is certainly not the case: The number of possible goals for a human agent is almost infinite. Here, we define desire as a computational process that suggests new goals to the planning system, and thus this is where the goals come from. (More details on how the desire system could actually choose goals will be given in Chapter 7, which also considers a different aspect of desire related to its interrupting and irresistible quality.)

A closely related concept is *aversion*, which is in a sense the opposite of desire. However, from a mathematical viewpoint, aversion is very similar to desire: The agent wants to *avoid* a certain state (or states) and wants to be in some other state. For example, the agent wants to be in a state in which some unpleasant object is not present. Thus, it is really a case of wanting and desire, just framed in a more negative way. I do not use the term aversion very much in this book since it is mathematically contained in the concept of desire. Whenever I use the word “desire”, aversion is understood to be included.<sup>7</sup>

## Intention as commitment to a goal

We have seen that a desire is something that suggests the goal of the agent. Note that I’m not saying that desire *sets* the goal, but it *suggests* a goal to the planning system. This difference is important because there might be conflicting goals; you don’t grab the chocolate every time you have desire for it. The agent needs to choose between different possible objects of desire. This is particularly important because attaining the desired goal state often takes time: The whole plan has to be executed from beginning till end, and new temptations—activations of the desire system which suggests new states as possible goals—may arise meanwhile. Some method of arbitrating between different desires is necessary.

Suppose the desired state for a monkey is one where the monkey has eaten a banana. The banana is currently high up the tree which is in front of the monkey, so the monkey needs to perform a series of actions to reach that desired state: it must climb up the tree, take the banana, peel it, and eat it. The monkey must thus figure out the right sequence of actions to reach the desired state—this is just the planning problem discussed above—and launch its execution.

But, suppose the monkey suddenly notices another banana in another tree near-by. Its desire system may suggest that the new banana looks like an interesting goal. The monkey now faces a new problem: Continue with the current banana plan, or set the new banana as a new goal? It may be common sense that after the monkey has launched the first banana plan, the monkey should, in most cases, persist with that plan until the end. The monkey should not start pondering, halfway up the first tree, whether it actually prefers to get the other banana in the other tree, even if it looks a bit sweeter. The key idea here is *commitment* to the current plan, and thus to a specific goal.

The reason why commitment is important comes fundamentally from computational considerations. Since computing a plan for a given goal takes a lot of computational resources, it would be wasteful to abandon it

---

<sup>7</sup>A linguistic confusion is created in English and many other languages in which it is commonplace to say “I don’t want X”, where X might be drilling noise in your office, or flies in your bedroom. What this actually means is that you want those things to be *absent*. It does *not* simply mean that you merely *refrain* from wanting that noise or the flies. You want “not X”, the opposite or absence of X, which is in fact the meaning of aversion.

too easily in favour of a new goal. It might be wasteful even to just consider alternative goals seriously, because that would entail computing possible plans to achieve all those goals. The agent must settle on one goal and one plan and execute it without spending energy thinking about competing goals.<sup>8</sup>

Another utility of commitment is that the agent has a better idea of what will happen in the future, and it can start planning further actions, e.g. a plan on what to do after reaching the goal of the current plan. So, while the monkey is climbing up the tree, it is a good idea to start thinking about the best way of getting the banana in the other tree *after* having grabbed the first banana. That would be planning the long-term future after completing the execution of the current plan; perhaps the monkey can directly jump to the other tree from the location of the first banana. Such long-term plans would obviously collapse if the monkey didn't first get the first banana due to lack of commitment, being distracted by yet another thing.

Commitment to a goal is also called *intention* in AI, and leads to an influential AI framework called belief-desire-intention (BDI) theory. Belief refers to the results of perception, which give rise to desires. BDI theory argues for the importance of intentions, as commitment to specific goals, on top of beliefs and desires.<sup>9</sup> Of course, there must be some limits to such commitment: If something unexpected happens, the goal may need to be changed. If a tiger appears, the monkey cannot persist with the goal of just eating a banana. Chapter 8 will consider the importance of emotions such as the fear aroused by the tiger as one computational solution.<sup>10</sup>

The concept of intentions has important implications for suffering, as will be discussed in detail in later chapters. To put it simply, I will propose that frustration and suffering are stronger if an intention is frustrated, as opposed to frustration of a simple desire as in the basic definition.

## Heuristics can help in planning

Still, we have not yet solved the central problem regarding the planning system. We saw above that because of the huge number of possible paths in planning, a complete tree search is quite impossible in most cases. Is planning then impossible? Fortunately, there are a couple of tricks and approximations that can be used to find reasonable solutions to the planning problem. Here we first consider what is called *heuristics*, while a more sophisticated solution is given in the next two chapters. These solutions also have important implications for the definition of frustration, and understanding suffering.

<sup>8</sup>A lot of physical energy would also be wasted if the monkey is already half-way up the tree and then decides to go for the other banana. But arguably that waste of energy would be taken into account by the monkey in its planning, so it does not need to be evoked as a separate reason for commitment. I think we can take here a viewpoint considering purely computational resources: Even if the monkey is intelligent enough to eventually understand this waste of physical energy after some thinking, it would still spend a lot of time and computational resources to reach that conclusion if there were no commitment mechanisms.

<sup>9</sup>(Bratman, 1987; Cohen and Levesque, 1990; Rao and Georgeff, 1991). See (Mulder, 2018; Brodaric and Neuhaus, 2020) for recent work and slightly different formulations. For a modern neuroscientific approach see (O'Reilly et al., 2014) which proposes something very similar using its concepts of "goal engagement", and "active goal". Note that the word "intention" has different meanings in the literature, and in particular this definition of the word is quite different from the meaning typically associated with "intentionality" *à la* Brentano. On the other hand, in the literature, there is some ambiguity on whether intentions are commitments to desires, goals, or plans. I consider them as commitments to goals.

<sup>10</sup>Some AI systems solve this problem by planning everything from scratch at regular intervals but that is unlikely to be possible in a real-time environment where the time needed for planning is the main bottleneck. In fact, my treatment here may not do full justice to Bratman's original definition of intention, where a plan is actually composed of several intentions. Such a definition creates more flexibility for behaviour in the sense that even if the circumstances change (or the circumstances were unpredictable to start with), the behaviour may flexibly move from one path to another by triggering an alternative sequence of intentions. The definition I use here is more similar to the later AI developments of the concept cited above.

A heuristic means some kind of method for evaluating each state in the search tree, usually by giving a number that approximately quantifies how good it is, i.e. how close to the goal it is. The point is that a heuristic does not need to be exact—if it were, we would have already solved the problem. It just gives a useful estimate, or at least an educated guess, of how “good” a state is.<sup>11</sup>

Sometimes, it is quite simple to program some heuristics in an AI agent. Consider a robot whose goal is to get some orange juice from the fridge and deliver it to its human master. Clearly, when the robot has orange juice in its hand, it is rather close to the goal; we could express that by a numerical value of, say, 8. If it is, in addition, close to its master, it is very close to its goal, say a value of 9. The most important thing is, however, to assist the robot at the beginning of the search, and that is where the heuristic is the most powerful. So, we could say that when the robot is close to the fridge, the heuristic gives a value of 2. When it has opened the fridge, the value is 3, and so on.

With such heuristics, the search task would not require that much computation. The robot just has to figure out how to get to some easily reachable state with a higher heuristic than the current state. Assuming the robot starts at an initial state with heuristic value 0, it would quickly compute that what it can achieve rather easily is a state of heuristic value of 2, by going to the fridge. The length of the tree to be searched for is thus much shorter, i.e. much fewer actions steps need to be taken in that subproblem. Once there, it only has to figure out how to open the door to get to the state with heuristic value of 3. Thus, the heuristic essentially divides a long complex search task into smaller parts. Each of these parts is quite short, so the exponential growth of the number of branches is much less severe.<sup>12</sup>

There is one famous success of AI where such tree search with heuristics was hugely successful: The Deep Blue chess-playing machine,<sup>13</sup> which beat the chess world champion, for the first time, in 1997. Its main strength was the huge number of sequences of moves (i.e. paths in a search tree) it was able to consider, largely because it was based on purpose-built, highly parallel hardware that was particularly good in such search computations on the chessboard. But its success was also due to clever heuristics, the main one being called “piece placement”, computed as the sum of the predetermined piece values with adjustments for location, telling how good a certain position is. (In chess, the state is the configuration of all the pieces on the board, and called a “position” in their jargon.)

Evolution has also programmed a multitude of heuristics in animals. Think about the smell of cheese for a rat. The stronger the smell, the closer the rat is to the cheese. The rat just needs to maximize the smell, as it were, and it will find the cheese. No complex planning is needed—unless there are obstacles in the way.<sup>14</sup>

However, the crucial problem is how to find such heuristics for a given planning problem. In fact, this is a very difficult problem, and there is no general method for designing them. Nevertheless, there is a general principle which has been found tremendously useful in modern AI, and can be used here as well: *learning*.

<sup>11</sup>A very general definition of a heuristic, not only applicable to the tree search problem, is given by Gigerenzer and Gaissmaier (2011): A heuristic is a strategy that ignores part of the information, with the goal of making decisions more quickly, frugally, and/or accurately than more complex methods.

<sup>12</sup>As a simple (and only approximative) numerical example, think of dividing a tree of length 20 into two parts. Each part of length 10 has  $1,024 = 2^{10}$  states, so the two search trees have total of 2,048 states. This is much less than the original tree with  $2^{20} = 1,048,576$  states.

<sup>13</sup>(Campbell et al., 2002)

<sup>14</sup>In fact, we see here that there is some intricate connection with heuristics and desires. When the rat smells the cheese, surely a desire for cheese appears in its system. See Chapter 8, and in particular footnote 24, on how the same computations can sometimes be interpreted as heuristics or desires.

Modern AI is very much about using learning from data as an approach to solving the problem of programming intelligence. In the case of planning, it turns out that a general approach for solving the planning problem is to learn to rate the states, i.e. learn to associate some kind of heuristic to each world state. This is why in the next two chapters, we delve into the theory of machine learning. Its specific application to solving the planning problem will be considered in Chapter 5, where we also consider a different approach to defining frustration.

## Chapter 4

# Machine learning as minimization of errors

In this chapter, we will go through some of the basics of the backbone of modern AI: machine learning. Such AI crucially relies on learning from incoming data, which is also true of the brain. Machine learning is most often used in conjunction with neural networks, which are powerful function approximators, loosely mimicking how computations happen in the brain. We will also consider an alternative, older approach to intelligence based on symbols, logic, and language, which is now called “good old-fashioned AI”. (The preceding chapter with its discrete, finite states, was an example of this latter approach.)

A central message in this chapter is that learning is often based on some measure of error. Minimizing such errors means optimizing the performance of the system. The fundamental importance of computing and signalling such errors is important in future chapters where such errors are directly linked to suffering, generalizing the concept of frustration. I conclude this chapter by claiming that any kind of learning from complex data can lead to quite unexpected results, something that the programmer could not anticipate.

### Neurons and neural networks

Modern AI is based on the observation that the human brain is the only “device” we know to be intelligent for sure and without any controversy. It is actually not easy to define what “intelligence” means, and I will not attempt to do that in this book.<sup>1</sup> Yet, nobody denies that the brain is intelligent—or, to put it another way, it enables us to behave in an intelligent way. The brain is intelligent as if by definition; it is the very standard-bearer of intelligence. If you want to build an intelligent machine, it makes sense to try to mimick the processing taking place in the brain.

#### *Neurons as tiny processors*

The computation in the brain is done by specialized cells called neural cells or neurons. A schematic picture of a neuron is in Fig. 4.1. A neuron receives input from other neurons, processes that input, and outputs the results of its computations to many other neurons. There are tens of billions of neurons in the human brain. Each single neuron can be seen as a simple information-processing unit, or a processor.<sup>2</sup>

<sup>1</sup>For standard textbook expositions on the definition of (artificial) intelligence, see e.g. Russell and Norvig (2020); Poole and Mackworth (2010). For particular viewpoints relevant to our discussions later, see e.g. Brooks (1991); Legg and Hutter (2007).

<sup>2</sup>It has also been claimed that other types of cells in the brain could also participate in computations, in particular glial cells (Perea et al., 2009). However, AI systems only use a single kind of cells, mimicking neurons.— I do not attempt to define “information pro-



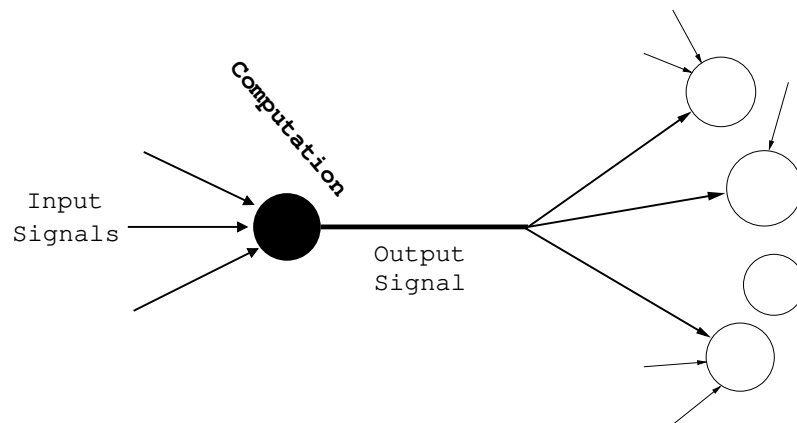


Figure 4.1: A schematic of a neuron. Input signals coming from other neurons (from the left) are received by the neuron (depicted by the black disk). Computation happens inside the neuron, and the resulting output signal is transmitted to a number of other neurons (depicted by white disks) on the right-hand side. The other neurons simultaneously receive input signals from many further neurons outside of this figure (depicted by further arrows).

All these tiny processors do their computations simultaneously, which is called *parallel* processing in technical jargon. The opposite of parallel processing is serial processing, where a single processor does various computational operations one after another—this is how ordinary CPU’s in computers work. Another major difference between the brain and ordinary computers is that processing in neurons is also *distributed*. This means that each neuron processes information quite separately from the others: It gets its own input and sends its own output to other neurons, without sharing any memory or similar resources. Compared to an ordinary PC, the brain is thus a *massively parallel and distributed* computer. Instead of a couple of highly sophisticated and powerful processors as found in a PC, the brain has a massive amount—billions—of very simple processors. (Parallel and distributed processing is discussed in detail in Chapter 6.)

While the actual neurons are surprisingly complex, in AI, a highly simplified model of a real neuron is used. Sometimes, such a model is called an artificial neuron to distinguish it from the real thing, but for simplicity, we call them just neurons. Like a real neuron, an artificial neuron gets input signals from other neurons, but each such input signal is very simple, just a single number; we can think of it as being between zero and one, like a percentage. Based on those inputs, the neuron computes its output which is, again, a single number. This output is, in its turn, input to many other neurons.

In such a simple model, the essential thing is to devise a simple mathematical formula for computing the output of the cell as a function of the inputs. In typical models, the output is computed, essentially, as a *weighted sum* of the inputs. The weights used in that sum are interpreted, in the biological analogy, as the strengths of the connections between neurons, or the incoming “wires” on the left-hand-side of Fig. 4.1. These weights can get either positive or negative values: The weight is defined as zero for those neurons from which no input is received. The weighted sum is usually further thresholded (i.e. passed through a nonlinear function) so that the output is forced to be between zero and one. In the brain, the connections are implemented through

---

cessing” in any rigorous way in this book. It is a very general concept with many meanings, and attempting to define it in a way that is both general and rigorous enough seems hopeless to me.

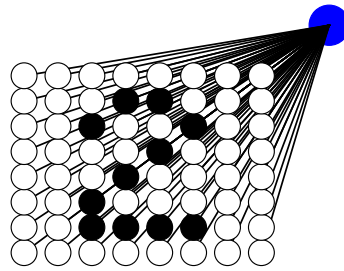


Figure 4.2: Synaptic weights of a neuron illustrated. Pixels shown in black have a connection strength of  $-1$  to the neuron (depicted in blue), while pixels shown in white have a connection strength of  $+1$ . The neuron is maximally activated when the input corresponds to the stored pattern, which is the digit “2”.

small communication channels called synapses, which is why the weights can also be called “synaptic”.

Importantly, these weights can be interpreted as a pattern, or a template, which the neuron is sensitive to. Thus, a neuron can be seen as a very simple *pattern-matching unit*. The neuron gives a large output if the pattern of all the input signals matches the pattern stored in the vector of weights or connection strengths.

As an illustration, consider a neuron which has a weight with the numerical value  $+1$  for inputs from another neuron, let’s call it neuron A, as well as a zero weight from neuron B, and a weight of  $-1$  neuron C. This neuron will give a maximal output (it is maximally “activated”) when the input to the neuron is similar to the pattern of those weights: It will output strongly when neuron A gives a large output and the neuron C has a small output, while it does not care what the output of neuron B might be. In other words, the neuron computes how well the pattern stored in its synaptic weights matches with the pattern of incoming input, and its output is simply a measure of that match.

Such pattern-matching is obviously most useful in processing sensory input, such as images. Consider a neuron whose inputs come from single pixels in an image. That is, the input consists of the numerical values of each pixel, telling how bright it is, i.e. whether it is white, black, or some sort of grey. Then, we can plot the synaptic weights as an image, so that the grey-scale value in each pixel in this plot is given by the corresponding synaptic weights. If they are  $-1$  or  $+1$  as in the previous example, we can plot those values as black and white, respectively. A neuron could have synaptic weights as in Fig 4.2. Clearly, this neuron is specialized for detecting a digit, in particular number two.

Of course, in reality, to recognize digits (or anything else) in real images, things are much more complicated. For one thing, the pattern to be recognized could be in a different location. If the digit is moved just one pixel to the right or to the left, the pattern matching does not work anymore, and the neuron will not recognize the digit. Likewise, if the digit were white on a black background instead of black on a white background, the same pattern-matching would not work. To solve these problems, we need something more sophisticated.

### *Networks based on successive pattern matching*

Building a *neural network* greatly enhances the capabilities of such an AI, and solves the problems just mentioned. A neural network is literally a network consisting of many neurons. Networks can take many different forms, but the most typical one is a hierarchical one, where neurons are organized into layers, each of which contains several cells, actually quite a few sometimes. The incoming input first goes to the cells in the first layer which compute their outputs and send them to neurons in the second layer, and so on. This is illustrated in Figure 4.3.

From the viewpoint of pattern-matching, such a network performs successive and parallel pattern matching. The input is first matched to all the patterns stored in the first-layer neurons, and those neurons then output the degrees to which the input matched their stored patterns or templates. These outputs are sent to the next layer, whose neurons then compare the pattern of first-layer activities to their templates. So, the second-layer patterns are not patterns of original input (such as the pixels of an image) but patterns of the first-layer activities, which form a description of the input on a slightly more abstract level. This goes on layer by layer, so that each neuron in each layer is “looking for” a particular kind of pattern in the activities of the neurons in the previous layer. The patterns are always stored in the synaptic weights of the neurons.

The utility of such a network structure is that it enables much more powerful computation. For example, consider the problem of a digit which could be in slightly different locations in the image, as mentioned above. The problem of different locations can be fixed by having several neurons in the first layer, each of which matches the digit in one possible location. All we need in the second layer is a neuron that adds the inputs of all first-layer neurons, and thus computes if any of them finds a match. With such a scheme, the second-layer neuron is able to see if there is a digit “2” at any location in the image.

## **Finding the right function by learning**

Now, a crucial question is how the synaptic connection weights can be set to useful values. In modern AI, the synaptic weights between neurons are *learned from data*, hence the term *machine learning*. Learning is really the core principle in most modern AI. Especially in the case of neural networks, it is actually difficult to imagine any alternative. How could a human programmer possibly understand what kind of strengths are needed between the different neurons? In some cases, it might be possible: in image processing, the first one or two layers do have rather simple intuitive interpretations, as we have alluded to above. However, with many layers—and neural networks can have thousands of them—the task seems quite impossible, and hardly anybody has seriously tried to design such neural networks by fixing the weights manually, based either on some theory or intuition.

In the brain, the situation is quite similar. There is simply not enough information in the genome—which is somewhat analogous to the programmer here—to specify what the synaptic connection strengths should be for all the neurons. It would hardly be optimal anyway to let the genes completely determine the synaptic connections, since animals live in environments that may change from one generation to another, and some individual adaptation to circumstances is clearly useful. What happens instead is that the synaptic weights change as a function of the input and the output of the neuron, or as a function of perceptions and actions of the organism. The capability of the brain to undergo such changes is called “plasticity”, and those changes are the biophysical mechanism underlying most of learning in humans or animals.

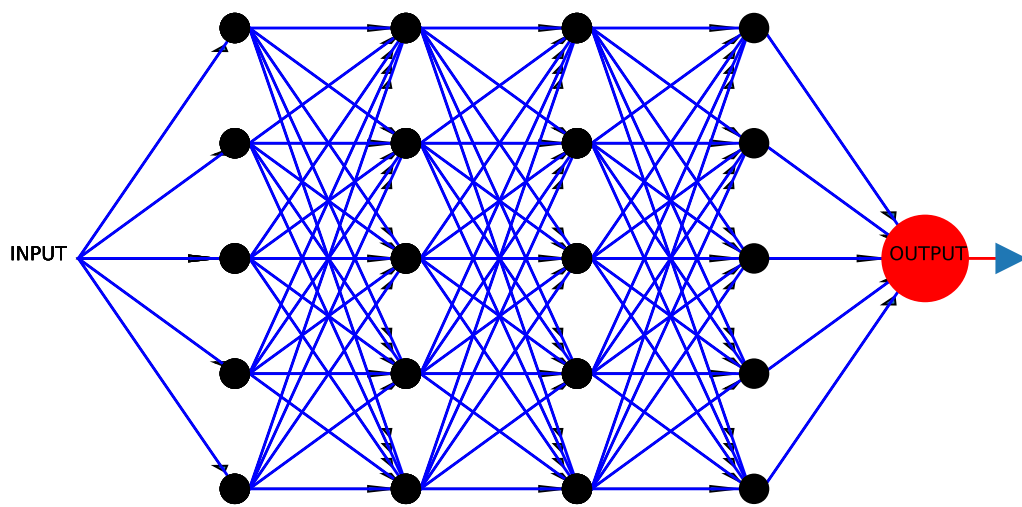


Figure 4.3: An illustration of a neural network. The information enters the system in the first “layer” of black neurons on the left-hand side. It is processed by several successive layers, each having five neurons illustrated by small black disks. Each neuron is doing a simple pattern-matching computation on its inputs, and transmitting the result of that computation to the next layer to its right, along the wires depicted in blue. The information is transmitted from the left (input) to the right (output). As a result of many neurons (in reality, thousands or even millions), the total computation of the network is highly complex and can achieve sophisticated object recognition, as well as many other kinds of computations.

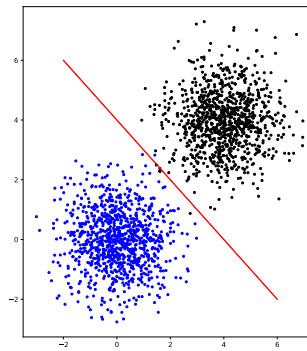


Figure 4.4: A simple illustration of what kind of a function a single neuron can learn in the basic case of classification with two classes. Each object (e.g. image of an animal) can be considered as a point in a very high-dimensional space where the coordinates correspond to pixel values, for example. For the purposes of this illustration, we assume there are only two input variables, so we can plot the points on a 2D plane. We also assume there are only two classes (something like “cats” and “dogs”) which correspond to black and blue points, respectively. In the ideal case, the neuron will learn to output a “one” when the input is in one of the classes, and a “zero” when it is in the other class. Such learning corresponds to learning the line that separates the two classes, drawn here as red. Finding a line that separates the classes is clearly possible based on this data, and you have probably done that automatically in your head while looking at this figure. Such learning can be done by a single artificial neuron due to the great simplicity of this illustration, but in reality, we would often need a neural network with many neurons and layers.

How such changes precisely happen in the brain is an immensely complex issue, and we understand only some basic mechanisms. Nevertheless, in AI, a number of relatively simple and very useful learning algorithms have been developed. Neural networks using them learn to perform basic “intelligent” tasks such as recognizing patterns (is it a cat or a dog?) or predicting the future (if I turn left at the next intersection, what will I see?). Learning in a neural network in such a case is based on learning a *mapping*, or function, from input data to output data. Let’s first consider a single neuron. It can basically learn to solve simple classification problems, as illustrated in Figure 4.4. If the classes are nicely separated in the input space, a single neuron can learn, as its synaptic weights, the pattern that precisely describes the difference between the classes.

However, a network with many neurons can learn to represent much more complex functions from input to output. The input data could be photographs and the output data could be a word describing the main content of the photograph (“cat”, “dog”, or “unicorn”). The learning of the input-output mapping then consists of changing the synaptic weights of all the neurons in all the layers. In successive layers, the network performs increasingly sophisticated and abstract computations, consisting of matching the inputs successively to the templates given by the weight vectors in each layer. After learning the right mapping, you can input a photograph to the system, and the output of the neural network will give its estimate of what the photograph depicts.

There is an infinite number of different ways you can use a neural network by just defining the inputs and outputs in different ways. If you want to learn to predict future stock prices, the input data would be the past

prices and the output, current stock prices. You can create a recommendation system that recommends new products to people in online shopping by defining the inputs to be some personal information of the customers, and the output whether the customer bought a certain item or not.<sup>3</sup> In a rather unsavoury application, the inputs would be what a social media user likes, and the output some sensitive personal information (say, sexual orientation), and then you can predict that sensitive information for anybody. Whether the prediction is accurate is another question, of course.

Here, we see one main limitation of machine learning: The availability of data. Where do you get the sensitive personal information of social media users in the first place, i.e. where do you get the data to train your network? Maybe nobody wants to give you such sensitive data. In other cases, the data may be very expensive to collect; for example, in a medical application, useful measurements and their analyses may cost a lot of money. Finding suitable data is a major limiting factor in neural network training; this is a theme we will come back to many times. Learning needs data, obviously; but it also needs the right kind of data, and enough of it.

## Learning as minimization of errors

After we have got the data, we need to define how to actually perform the learning. Most often, the learning is based on formulating some kind of error, and the network then tries to minimize it by an algorithm. The error is a function of the data, i.e. something that can be computed based on the data at our disposal, and tells us something about how well the system is performing.

Suppose the data we have consists of a large number of photographs and the associated categories (cat/dog etc.). To recognize patterns in the images, the network could learn by minimizing the percentage of input images classified incorrectly, called classification error. Alternatively, suppose we want to learn to predict how an agent's actions change the world—say, how activation of an artificial muscle changes the position of the arm of a robot. In that case, what is minimized is prediction error: The magnitude of the difference between the predicted result of the action and the true result of the action (which can be observed after the action). Such errors don't usually go to zero, i.e. there will always be some error left even after a lot of learning. This is due to the uncertain and uncontrollable nature of the world and an agent's actions; that's another theme we will discuss in detail in later chapters.

We also need to develop an algorithm to minimize the errors, but there are standard solutions that usually are satisfactory. What most such algorithms have in common is they learn by making tiny changes in the weights of the networks. This is because optimizing an error function, such as classification error, is actually a very difficult computational task: There is usually no formula available to compute the best values for the weight vectors. In contrast, what is usually possible is to obtain a mathematical formula that gives the *direction* in which the weights should be changed to make the error function decrease the fastest. That direction is given by what is called the *gradient*, which is a generalization of the derivative in basic calculus.

So, you can optimize the error function step by step as follows. You start by assigning some random values to the weight vectors. Given those values, you can compute the gradient, and then take a small step in that direction (i.e. move the weight vector a bit in that direction), which should reduce the error function, such as prediction error. But you have to repeat that many times, often thousands or even millions, always computing

---

<sup>3</sup>(Davidson et al., 2010)

the gradient for the new weight values obtained at the previous step. (The direction of the gradient is different at every step unless the error function is extremely simple.) Such an algorithm is called *iterative* (repeating) because it is based on repeating the same kind of operations many times, always feeding in the results of the previous computation step to the next step. Using an iterative algorithm may not sound like a very efficient way of learning, but usually an iterative gradient algorithm is the only thing we are able to design and program.

Such an algorithm is a bit like somebody giving you instructions when you're parking your car and cannot see the right spot precisely enough. They will only give instructions which are valid for a small displacement. When they say "Back", that means you need to back the car *a little bit*, and then follow some new instructions. This is essentially an iterative algorithm, where you get instructions for the direction of a small displacement, and they are different at every time step.

Neural networks use an even more strongly iterative method, based on computing the gradient for just a few data points at a time. A data point is one instance of the input-output relation data, for example, a single photograph and its category. In principle, a proper gradient method would look at all the data at its disposal, and push the weights a small step in the direction that improves the error function, say the classification error, for the data set as a whole. However, if we have a really big data set, say millions of images, it may be too slow to compute the error function for all of them. What the algorithms usually do is to take a small number of data points and compute the gradient only for those. That is, you just take a hundred photographs, say, and compute the gradient, i.e. in which direction the weights should be moved to make the classification accuracy better, for those particular images. Importantly, at every step you randomly select a new set of a hundred images, and do the same thing for those images.

The point is that you are still *on average* moving the weights in the right direction, so this is not much worse than computing the real gradient. But crucially, you can take steps much more quickly, since the computation of the gradient is much faster for the small sample. It turns out that in practice, the benefit of taking more steps often overwhelms the slight disadvantage of having just an approximation of the gradient.<sup>4</sup>

Putting these two ideas together, we get what is called the *stochastic gradient descent* algorithm. Here, "descent" refers to the fact that we want to minimize an error. "Stochastic" means "random", and refers to the fact that you are computing the gradient for randomly chosen data points, so you are going in the right direction only on average.

Suppose you're in an unfamiliar city and you need to get to the railway station. Your "error function" is the distance from the station. You can ask a passer-by which direction the station is, and you get something analogous to the gradient for one data point. Now, of course, that direction given by the passer-by is not certain, she could very well be mistaken; maybe she even said she is not quite sure about the direction. But you probably prefer to walk a bit in that direction, and then ask another passer-by. This is like stochastic gradient descent, where you follow an approximation of the gradient, given by each single data point. The opposite would be that instead of following each person's advice one after the other, you first ask everybody you see on the street where they think the station is, and move in the average of the directions they are giving. Sure, you would get a

---

<sup>4</sup>To this advantage we have to add the more technical one that stochastic methods include an implicit regularization and are thus less likely to overfit the data (Bottou, 2003; Hardt et al., 2016). Overfitting is an important problem in practical AI learning, but I don't discuss it at any length in this book. Basically, it means that if the amount of data at your disposal is limited (as it almost always is), the learning may go wrong in a particular way. The learning may seem to work well for the data you have, achieving a good "fit", but the predictions your neural network gives are actually useless, because the learning "overfit" your data and does not work (or give a good fit) for any new data on which you would like to apply the system in the future.

very precise idea of what the right direction is, but you would advance very slowly—this is analogous to using the full, non-stochastic gradient.

## Gradient optimization vs. evolution

There are many more ways of optimizing an error function, and many systems that can be conceptualized as the optimization of a function. In particular, evolution is a process where the error function called *fitness* is optimized. Fitness is basically the same as reproductive success, which can be quantified as the expected number of offspring of an organism.<sup>5</sup>

Fitness is, of course, maximized, while errors in AI are minimized. However, this difference is completely insignificant on the level of the optimization algorithms, since maximization of a function is the same as minimizing the negative of that function. Thus, evolution can equally well be seen as minimization of the negative fitness.

In general, such a function to be optimized—whether minimized or maximized—is called an “objective function”. The objective function does not necessarily have to be any kind of a measure of an error, although in AI, it often is. Note that the objective function is different from the function from the input to the output that the neural network is computing, as described above. The objective function is what enables the system to learn the best possible input-output function, so it works on a completely different level.

Evolution works in a very different way than stochastic gradient descent. But it is actually possible to mimic evolution in AI and use what is called evolution strategies, evolutionary algorithms, or genetic algorithms. These are iterative algorithms which are sometimes quite competitive with gradient methods. They can optimize any function, which does not need to have anything to do with biological fitness. For example, we can learn the weights in a neural network by such methods. The idea is to optimize the given error function by having a “population” of points in the weight space, which is like a population of individual organisms in evolution.

Like real evolution, such algorithms are based on two steps. First, new “offspring” is generated for each existing “organism”. In the simplest case, you randomly choose some new weight values close to the current weight values of each organism, which is a bit like asexual reproduction in bacteria, with some mutations to create variability. Then, you evaluate each of those new organisms by computing the value of the error (such as classification error) for their values for weights. Finally, you consider the value of the error as an analogue of fitness in biological evolution, albeit with the opposite sign because fitness is to be maximized while an error function is to be minimized. What this means is that you let those organisms (or weight values) with the smallest values of the error “survive”, i.e. you keep those weight values in memory and discard those weight

---

<sup>5</sup>While the idea of evolution as fitness maximization can be found in many textbooks, some biologists would refute the whole idea of evolution optimizing any single function; see a recent review by Birch (2016). To some extent, this controversy may also have arisen because of some semantic confusion about whether that would mean that evolution has already optimized the function (which would be a very strong statement) or whether it is *in the process* of optimizing it (which may be more plausible); see Parker and Smith (1990). Regarding the precise definition of fitness, it seems impossible to find a consensus opinion (Rosenberg and Bouchard, 2011; Grafen, 2008). A standard textbook definition would be along the lines “expected number of offspring”, but this may have to be complemented by the concept of inclusive fitness treated in Chapter 6. — In this book, I tend to anthropomorphize evolution rather unashamedly, often comparing it to a human programmer. I believe that is a useful pedagogical device since humans find it easier to think of natural phenomena in terms of agents that have goals instead of a more abstract description such as a dynamical system. See e.g. Dawkins (1986) for a rather strictly anti-anthropomorphic view on evolution.



values which have larger (that is, worse) values of the error function. You also discard the organisms of the previous iteration or “generation”, since those individual organisms die in the biological analogy.

Such an evolutionary algorithm will find new weight values which are increasingly better because only the organisms with the best weight values survive in each iteration. Thus, it is an iterative algorithm that optimizes the error function. It is a randomized algorithm, like stochastic gradient descent, in the sense that it randomly probes new points in the weight space. In fact, an evolutionary algorithm is much more random than stochastic gradient descent, since gradient methods use information about the shape of the error function to find the best direction to move to, while evolutionary methods have no such information. This is a disadvantage of evolutionary algorithms, but on the other hand, one step in an evolutionary algorithm can be much faster to compute since you don’t need to compute the gradient, just random variations of existing weights.<sup>6</sup>

So, we see that *both evolution and machine learning are optimizing objective functions*. The optimization algorithms are often quite different, but they need not be. One important difference is that in machine learning, the programmer knows the error function, and explicitly tells the agent to minimize it. In real biological evolution, fitness is an extremely complicated function of the environment; it cannot be computed by anybody, nor can its gradient. Real biological fitness can only be observed afterwards, by looking at who survived in the real environment, and even then you only get a rough idea of the values of fitness of the individual organisms concerned: Those who die *probably* had a low fitness, but it is all quite random—even more than stochastic gradient descent. (If an organism were actually able to compute the gradient of its fitness, that would give it a huge evolutionary advantage.) Another important difference is that in biology, evolution works on a very long time scale, over generations, while in AI, the learning in the neural networks happens typically inside an individual’s life span. The evolutionary algorithms in AI typically learn within an individual agent’s lifespan as well, only simulating “offspring” of a neural network in its processors.

## Learning associations by Hebbian rule

So far, we have seen learning as based on finding a good mapping from input to output. Such learning is called *supervised* because there is, metaphorically speaking, a “supervisor” that tells the network what the right output is for each input. Yet, sometimes it is not known what the output of a neural network should be, or whether there is any point at all in talking about separate input and output—especially if we are talking about the brain. In such a case, learning needs to be based on completely different principles, called *unsupervised learning*. In unsupervised learning, the learning system does not know anything about any desired output (such as the category of an input photo). Instead, it will try to learn some regularities in the input data.

The most basic form of unsupervised learning is learning associations between different input items. In a neural network, they are represented as connections between the neurons representing those two items. For example, if you have one neuron representing “dog” and another neuron representing “barking”, it is reasonable that there should be a strong association between them.

One theory of how such basic unsupervised learning happens in the brain is called Hebbian learning. Donald Hebb proposed in 1949 that when neuron A repeatedly and persistently takes part in activating neuron B,

---

<sup>6</sup>For a basic genetic algorithm, see for example (Such et al., 2017) and the references therein. A particular advantage of evolutionary methods is that often they can be very efficiently parallelized, even if combined with gradient methods by using a stochastic gradient descent to generate the offspring (Salimans et al., 2017); parallelization is explained in Chapter 11.

some growth process takes place in one or both neurons such that A's efficiency in activating B is increased.<sup>7</sup> Such Hebbian learning is fundamentally about learning associations between objects or events. A very simple expression of the Hebbian idea is that “cells that fire together, wire together”, where “firing” is a neurobiological expression for activation of a neuron. In this formulation, Hebbian learning is essentially analysing statistical correlations between the activities of different neurons.<sup>8</sup>

One thing which clearly has to be added to the original Hebbian mechanism is some kind of forgetting mechanism. It would be rather implausible that learning would only increase the connections between neurons. Surely, to compensate, there must be a mechanism for decreasing the strengths of some connections as well. Usually, it is assumed that if two cells are not activated together for some time, their connection is weakened, as a kind of negative version of Hebb's idea.<sup>9</sup>

Hebbian learning has been widely used in AI, and it has turned out to be a highly versatile tool. You can build many different kinds of Hebbian learning, depending on how the inputs are presented to the system and on the mathematical details of how much the synaptic strengths are changed as a function of the firing rates. You can also derive Hebbian learning as a stochastic gradient descent for some specially crafted error functions.<sup>10</sup>

## Logic and symbols as an alternative approach

The inspiration for neural networks is that they imitate the computations in the brain. Since the brain is capable of amazing things, that sounds like a good idea. But historically, before neural networks, the initial approach to AI was quite different. It was actually more like the world of planning we saw in Chapter 3, where the world states are discrete, and there are few if any continuous-valued numerical quantities.

In early AI, it was thought that logic is the very highest form of intelligence, and therefore, AI should be based on logic. Also, the principles of logic are well-known and clearly defined, based on hundreds of years of mathematics and philosophy, so they should provide, it was thought, a solid basis on which to build AI. In modern AI, such logic-based AI is not very widely used, but it is making a come-back: It is increasingly appreciated that intelligence is, at its best, a combination of neural networks and logic-based AI—now called “good old-fashioned” AI, or GOFAI for short. Such logic provides a form of intelligence that is in many ways completely different from neural network computations, as we will see next.

### *Binary logic vs continuous values*

Mathematical logic is based on manipulating statements which are connected by operators such as AND and OR. For example, a robot might be given information in the form of a statement that “the juice has orange

---

<sup>7</sup>Adapted from Hebb (1949), page 62. I have simplified the original quote, stripping it from its neurobiological terminology. The original formulation is “When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased.”

<sup>8</sup>This is perhaps oversimplifying the original idea: Recent research in neuroscience has emphasized that, as in the original definition above, it is important that cell A participates in the activation, i.e. is has a causal influence of cell B. This would usually mean that cell A is activated before cell B (Markram et al., 2012). However, in most implementations of Hebbian learning in AI, such causal and temporal aspects are not used. Actually, the details of how Hebbian learning works in the brain are not very well understood.

<sup>9</sup>(Oja, 1982; Zenke et al., 2017)

<sup>10</sup>(Oja, 1992; Hyvärinen and Oja, 1998)

colour AND the juice is in the fridge”. Any statement can also be made negative by the NOT operator. An important assumption in such systems, in their classical form, is that any statement is either true or false; no other alternatives are allowed. This goes back to Aristotle and is often called the law of the “excluded third”. That is, truth values are binary (have two values).

Such logic is perfectly in line with the basic architecture of a typical computer. Current computers operate on just zeros and ones, and those zeros and ones can be interpreted as truth values: Zero is false and one is true. Such computers are also called “digital”, meaning that they process only a limited number of values, in this case just two. Our basic planning system in the preceding chapter, with its finite number of states of the world, was an example of building AI with such a discrete approach, and planning is fundamentally based on logical operations.

The brain, in contrast, computes with quantities which are in “analog” form, which means the continuous-valued numbers, which take a potentially infinite number of possible values. Artificial neural networks do exactly the same, as they are trying to mimic even this aspect of the brain. It is rather unnatural for the brain to manipulate binary data or to perform logical operations. It is possible only due to some very complex processes which we do not completely understand at the moment.

This distinction between digital and analog information-processing is another important difference between ordinary computers on the one hand, and the real brain or its imitation by neural networks on the other. (Earlier we saw the distinction between parallel and distributed processing in the brain versus the serial processing in an ordinary computer.) The digital nature of ordinary computers implies that any data that you input has to be converted to zeros and ones. This is actually a bit of a problem because a lot of data in the real world does not really consist of zeros and ones. For example, images are really intensities of light at different wavelengths, measured in a physical unit called “lux”. One pixel in an image might have an intensity of 1,536 lux and another 5,846 lux. It is, again, rather unnatural to represent such numbers using bits, which is why processing non-binary data such as images is relatively slow in modern computers, compared to binary operations.<sup>11</sup>

### *Categories and symbols*

Saying that things are either true or false is related to thinking in terms of *categories*. Human thinking is largely based on using categories: We divide all the perceptual input—things that we see, hear, etc.—into classes with little overlap. Say, you divide all the animals in your world into categories such as cats, dogs, tigers, elephants, and so on, so that each animal belongs to one category—and usually just one. Then, you can start talking about the animals in terms of true and false. You can make a statement such as “Scooby is a dog”, and that is either true or false based on whether you included that particular animal in the dog category; any other (third) option is excluded.

Categories are usually referred to by *symbols*, which in AI are the equivalent of words in a human language. For example, we have a category referred to by the word “cat”, which includes certain “animals” (that’s another category, actually, but on a different level). Ideally, we have a single word that precisely corresponds to each

---

<sup>11</sup> I shall just briefly mention another crucial difference between brains and ordinary computers: An ordinary computer has hardware and software, and these two are separate. The same hardware can run different kinds of software, and the same software can be used on different hardwares. In fact, you can take software from one computer and download it to another, similar computer and it will work on that new computer as well. However, in the brain, it is difficult to see any clear distinction between the software and the hardware: nobody ever downloaded software into their brain. Such a division between hardware and software is part of what is called the von Neumann architecture, named after the great mathematician John von Neumann.

single category, like the words “cat” and “animal” above. Such symbols are obviously quite arbitrary since in different languages the words are quite different for the same category. An AI system might actually just use a number to denote each category.

We see that logic-based processing goes hand-in-hand with using categories, which in its turn leads to what is sometimes called symbolic AI. These are all different aspects of GOF AI.

#### *From hand-coded logic to learning*

Historically, one promise of GOF AI was to help in medical diagnosis, where the programs were often called “expert systems”. This sounds like a case where categories must be useful since medical science uses various categories referring to symptoms (“cough”, “lower back pain”) as well as diagnoses (“flu”, “slipped disk”).

The basic approach was that a programmer asks a medical expert how a medical diagnosis is made, and then simply writes a program that performs the same diagnosis, or makes the same “decisions” in the technical jargon. For example, one decision-making process by the human expert might be translated into a formula such as

IF cough AND nasal congestion AND NOT high fever THEN diagnosis is common cold

However, this research line soon ran into major trouble. The main problem was that medical doctors, and indeed most human experts in any domain, are not able to verbally express the rules they use for decision-making with enough precision. This is rather surprising since we are operating with human language and well-known categories. The situation is different from neural networks where it is intuitively clear that no expert can directly tell what the synaptic weights should be, because their workings are so complex and counterintuitive. Yet, it turned out that even medical diagnoses are often based on intuitive recognition of patterns in the data, which is a form of *tacit knowledge*. Tacit knowledge means knowledge, or skills, which cannot be verbally expressed and communicated to others.

A major advance in such early AI was to understand that expert systems should actually learn the decision rules based on data. Again, learning provides a route to intelligence that is more feasible than trying to directly program an intelligent system. Given a database with symptoms of patients together with their diagnoses given by human experts, a machine learning system can learn to make diagnoses. Such learning is not so fundamentally different from learning by neural networks. What is different is that the data is categorical (“cough”, “no cough”), and the functions are computed in a different way, for example by combining logical operations such as AND, OR, and NOT.

#### *Categorization and neural networks*

By definition, such logic-based AI can only learn to deal with data which is given as well-defined categories. Yet, real data is often given as numbers instead of categories; even medical input variables often include numerical data in the form of lab test results. In this medical diagnosis, we have indeed a category called “high fever”. The system is effectively dividing the set of possible body temperatures into at least two categories, one of which is “high fever”. How are such categories to be defined? What is fever? What is low fever and what is high fever? Here we see a deep problem concerning how categories should be defined based on numerical data, such as sensory inputs.

Again, some progress can be made by learning, this time learning the categories themselves from data. The AI can consider a huge number of possible categorizations of body temperatures: It could try setting the threshold for high fever at any possible value. If there is enough data on previous diagnoses by human doctors, the system could use that to learn the best threshold. In fact, the right threshold could be found as the one that minimizes classification error.

However, while it is possible to learn categories in such very simple numerical data, GOF AI has great difficulties in processing complex numerical data. It is virtually impossible to use it to process high-dimensional sensory input, such as images consisting of millions of pixels. This was a rather big surprise for GOF AI researchers in the 1970s and 80s. After all, categorization of visual input is done so effortlessly by the human brain that it may seem to be easy. Yet, AI researchers working in the GOF AI paradigm found it to be next to impossible. The early research on GOF AI was fundamentally over-ambitious, grossly underestimating the complexity of the world, as well as the complexity of the brain processes we use to perceive and make decisions.

One reason for the current popularity of neural networks is that processing high-dimensional sensory data is precisely what they are good at. In fact, it is clear by now that neural networks operate in a completely different regime from such logic-based expert systems. There are no categories, and no symbols, in the inner workings of neural networks: What they typically operate on is numerical, sensory input such as images, or some transformations of such sensory input. Raw grey-scale values of pixels are kind of the opposite of neat, well-defined categories.

As such, neural networks and logic-based systems can complement each other in many ways. Usually the categories used by a logic-based system need to be recognized from sensory input, so a neural network can tell the logic-based system whether the input is a cat or a dog. In particular, a neural network can take sensory data as input, and its output can *identify the states* used in action selection; to begin with, it can tell the planning system what the current state of the agent is. In Chapter 7 we will consider in more detail this fundamental distinction between two different modes of intelligent information-processing, which are found both in AI and human neuroscience.

## Emergence of unexpected behaviour

Finally, let me mention a phenomenon that is typical of any learning system. “Emergence” means that a new kind of phenomenon appears in some system due to complex interactions between its parts. It is a special case of the old idea, going back to Aristotle at least, that “the whole is more than the sum of its parts”. For example, systems of atoms have properties that atoms themselves do not—consider the fact that a brain can process information while single atoms hardly can. Likewise, evolution is based on emergence. Its objective function is given by evolutionary fitness, which sounds like a very simple objective function. Nevertheless, it has given rise to enormous complexity in the biological world, as well as human society. What is typical of such emergence is that its result is extremely difficult to predict based on knowledge of the laws governing the system. If the objective function given by fitness had been described to some super-intelligent alien race a few billion years ago, they would hardly have been able to predict what the world looks like these days.

Machine learning is really all about the emergence of artificial intelligence. We build a simple learning algorithm and give it a lot of data, and hope that intelligence emerges. It is the interaction between the algorithm and the data that gives rise to intelligence. This seems to work, if the algorithm is well designed, there is enough data of good quality, and sufficient computational power is available. Such emergence in machine learning is

actually a bit different from emergence in other scientific disciplines. In physics, very simple natural laws *by themselves* can give rise to highly complex behaviour. In machine learning, the complexity of the behaviour of the system is, to a large extent, a function of the complexity of the *data*. In some sense, one could even say the complex behaviour learned by an AI does not emerge but is extracted, or “distilled”, from input data. The complexity of the input data is due to the complexity of the real world, which is obvious when inputting a million photographs into a neural network.

The emergent nature of the behaviour learned by an AI implies that, just like in evolution, there is often something unexpected in the resulting system. The complexity of the input data usually exceeds the intellectual capacities of the programmer. So, the programmer of an AI cannot really know what kind of behaviour will emerge: Often the system will end up doing something surprising.

In this book, we will encounter several forms of emergent properties in learning systems which are related to suffering. While some kind of suffering may be necessary as a signal that things are going wrong, we will also see how an intelligent, learning system will actually undergo *much more* suffering than one might have expected. To put it bluntly, a particularly intelligent system will find many more errors in its actions and its information-processing. In fact, finding such errors was necessary to make it so intelligent in the first place. Therefore, a learning system may learn to suffer much of the time, even though that is not what the programmer intended.

## Chapter 5

# Frustration due to reward prediction error

Now, armed with modern machine learning theory, we revisit the problem of action selection and the concept of frustration. In planning treated in Chapter 3, the main computational problem is looking several steps ahead, which can lead to quite impossible demands of computational capacities. Another constraint is that it requires a model of how your actions affect the world, i.e. where do you go in the search tree when you perform a given action in a given state. As such, planning is not really a good method for action selection if computational resources are very limited, as in a simple computer, or a very simple animal such as an insect.

In this chapter, we consider an alternative way to action selection, based on learning. A paradigm called reinforcement learning enables learning intelligent actions without any explicit planning, thus avoiding many of its problems. It also generalizes the framework of a single goal to maximization of rewards obtained at different states. While it can be performed even in very simple animals and computers, it is also used by humans; it is similar to how habits work.

We then consider how frustration can be defined in such a case; it can no longer be simply defined as not reaching the goal—since there is no explicit goal. We define more general error signals called reward loss and reward prediction error, which have been linked to signals of certain neurons in the mammalian brain. Thus, we expand the view where frustration is related to error signalling by linking it to errors in prediction.

Repeated frustration is thus something necessary for learning algorithms to work, and intelligence may not be possible without some frustration. We further see how the very construction of an agent based on reward maximization means that it is insatiable, never satisfied with the amount of reward obtained. Moreover, it can be directed towards intermediate goals which are not valuable in themselves, but simply predictive of future reward. Evolutionary rewards, in particular, can lead to behaviour which resemble obsessions.

## Maximizing rewards instead of reaching goals

In modern AI, action selection is most often not based on planning, but a framework where the obtained *rewards*, or reinforcement, is maximized. This is useful because often an AI does not have just a single goal to accomplish, but many things it should take care of. Defining behaviour as maximization of rewards as opposed to reaching goals is also often thought to be more appropriate for modelling behaviour in simple animals, which are thought to be incapable of the sophisticated computations needed in planning (more on this in Chapter 9).

For example, if a cleaning robot disposes of some dust in the dustbin, it could be given a reward signal. Since there are many rooms and many dustbins in the building, it makes sense to give a reward whenever the

robot disposes of some of the dust. In principle, we could decide to give it a single reward when all the rooms are completely clean; however, it is common sense to rather give it a reward every time it removes some dirt or dust from any of the rooms. After all, the robot has done something useful every time it reduces the amount of dust in the room; telling this to the robot is highly useful information, and it would simply complicate the learning if the reward were postponed until the robot has completed some larger part of the task.

In fact, giving a single reward at the end would mean the robot has to engage in long-term planning, which is difficult. A “piece-wise” training by giving rewards for small accomplishments is not very different from how you would teach a child to perform a rather demanding, long task, say tying shoelaces: Divide the task into successive parts and give the child a small encouragement when it completes each small part. This is computationally advantageous since it eliminates the need for long-term planning, a bit like the heuristics we saw above.<sup>1</sup>

Reinforcement or reward can also be negative; if the robot tries to put household items in the dustbin, it can be given some. Negative reinforcement is really what we usually call a punishment—but the word is interpreted without any moral connotations here.

Thus, we actually ground action selection in the optimization of an objective function, i.e. a quantity to be optimized. Earlier, we saw that minimization of an error function, such as the number of images incorrectly classified, is the way an AI can learn to recognize objects in images. Here we define a different kind of objective function which is the basis of action selection: It is equal to the *sum of all future rewards*. It is a function of the action selection parameters of the agent, and more precisely, it expresses how much reward the agent can obtain by behaving according to its current action selection system.

Such a learning process based on maximization of future rewards by learning a value function is called *reinforcement learning*. Reinforcement learning can be seen as a third major type of learning in AI, in addition to supervised and unsupervised learning.

In a sense, this future reward is the ultimate objective function of an agent. Its maximization, by tuning the action-selection system, is the very meaning of life of the agent. The objective functions we saw earlier, used to learn things like pattern recognition by minimization of errors, are there merely to help in maximizing this reward-based objective function.

In such a reward-based objective function, more weight is often put on the rewards in the near future as opposed to rewards in the far-away future, which is called discounting. The justification for this is complicated, but suffice it to say that such discounting is often evident in human behaviour: Humans prefer to have their reward right now, and value it less if they have to wait. To keep the discussion simple, I sometimes ignore discounting in what follows, but it could be used in almost every case considered in this chapter.<sup>2</sup>

---

<sup>1</sup>It is also essential in training animals to perform long sequences of actions; in that context it is called “shaping” (Krueger and Dayan, 2009; Ng et al., 1999). However, reinforcement learning is a much more general concept than just dividing a long sequence into smaller parts. In the case of the cleaning robot, there may not be any end to the cleaning task since more dust appears constantly. The only meaningful goal for the cleaning robot may be to just remove dust and dirt as much as possible, which is exactly captured by the reward formalism.

<sup>2</sup>For discussions of different kinds of discounting, and in particular for comparisons between exponential and hyperbolic discounting, see Dasgupta and Maskin (2005); Ainslie (2001).



## Learning to plan using state-values and action-values

As such, the sum of future rewards gives a more general framework than having a single goal as in Chapter 3, since trying to reach a single goal can be accommodated in the reward framework by simply giving a reward when the agent reaches the goal, and no reward otherwise. In such a case, discounting further means the agent receives more reward if it reaches the goal more quickly, which is intuitively reasonable.

It turns out that we can use this reformulation of planning as reward maximization to our advantage, since the algorithms developed for maximizing future rewards give a particularly attractive way of solving the problem of planning. In Chapter 3, we saw how difficult planning is due to the exponential explosion in the number of possible plans to choose from. While heuristics were proposed as a practical trick to make the computations more manageable, there is no universal way of designing good heuristics.

Like in other branches of AI, it has been found that *learning* solves many of these problems. Intuitively, if the agent encounters the same planning problem again and again, it can store information about the previous solutions (or attempts) in memory. For example, a cleaning robot will probably clean the same building many, many times, and a delivery robot will deliver the parcels to the same addresses quite a few times. So, such agents should be able to learn something about planning in their respective worlds. This would be a clear improvement compared to heuristics which need to be explicitly programmed in the system by programmers, as in our examples above, and it is often not at all clear how to do that.<sup>3</sup>

In reinforcement learning, there is a sophisticated mathematical theory that tells how to learn a particularly good substitute of a heuristic, called the *state-value function*. It is a clever way of learning to deal with the complexity of the search in a planning tree. The basic principle is simple: Using the previous planning results in its memory, the agent can compute something like a heuristic based on how well it performed starting from each possible state. If it found the goal quickly starting from a certain state, that state gets a large state-value.

In the case where we have a single goal, the state-value function basically tells you how far from the goal you are, thanks to discounting which takes account of the time needed to reach the goal. A delivery robot that frequently delivers stuff to the same building (say town hall) would easily learn the distance from any other building to the town hall. In the beginning, when it had a delivery to the town hall, it had to spend a lot of time and effort in planning the path there. But little by little, it gained information by storing any results of executed plans in its memory, and learned the distance from any other building to the town hall. Such distances now give the state-value function for that goal (the state-value is actually a decreasing function of that distance). When the robot next needs to go to the town hall, it recalls the distances—to the town hall—from those buildings that are close to its current location, and simply decides to move in the direction of the near-by building which has the smallest distance to the town hall. Thus, it has learned a kind of a heuristic that avoids planning action sequences altogether.

This works even in a very general setting when there is no particular goal. In general, the value of a state is defined as the sum of all future rewards the agent can obtain *starting* from that state.<sup>4</sup> After successful learning

---

<sup>3</sup>Actually, Deep Blue mentioned in Chapter 3 did already use some learning as well: It analysed data from several chess databases, including 700,000 historical games played by human grandmasters, to compute another heuristic.

<sup>4</sup>Next, I give a more rigorous and general definition of state-values. To begin with, it must be noted that the state-value is a function of the “policy” used by the agent; the policy is what I call the action selection system in the text, i.e., a system deciding which action is selected (and with which probability) in any given state. Further, we have to take into account the fact that the world may have some randomness in it, so we have to consider expected reward. The value function at a given state is then generally defined as the expected amount of discounted reward that the agent will obtain starting from that state, when it follows that policy. (Sometimes,

of the state-values, the solution to the problem of action selection is that at each step, whatever state the agent may be in, the agent just selects the single action which leads to the state with the highest state-value (e.g. closest to the town hall above). There is no need to compute several steps ahead, or make a search within the huge search tree anymore. This reduces the complexity of the computations radically: instead of planning all its actions up to reaching the goal, the state-values now provide kind of intermediate goals, one step ahead, that the agent can very easily reach. However, a lot of time and computation still needs to be spent on first learning the state values.<sup>5</sup>

### *Completely reactive action selection by action-values*

While we have thus solved the problem of the computational explosion of planning, there is still the problem that the agent needs to have a model of how the world works. Even using the state-values, it needs to understand which action takes it into a state with higher state-value. Now, consider an extreme case where the agent has no model of how the world works in the sense that it has no idea what about the effects of its actions. Then, it is not enough to assign values to different states since the agent does not know how to get from state A to state B. (Still, we assume the agent knows in which state it is, at any given moment, so it does have some minimal model of the world.)

The trick to learning to act even with such a minimal model of the world is to learn what is called the *action-value* function. When the agent is in a given state, the action-value function tells the value of each of its actions, in terms of how much the total future reward is if the agent performs that action.<sup>6</sup> This makes action selection

---

when speaking about the state-value, it is more specifically assumed that the policy in question is the optimal policy which gives the highest expected reward.) This definition simplifies to the definition we just gave for the case of a single goal in a deterministic world, where the state-value is a decreasing function of the distance to the goal. The connection can be seen by defining that there is a reward at the goal and nowhere else, and using the fact that there is discounting, and thus rewards in the distant future are given less weight than rewards in the near future. Then, the closer you are to the goal, the larger the expected reward is, because the reward at the goal is given more weight when you are closer to the goal. (I define here “closer” to mean that you can get there more quickly compared to the situation where you are further away and need time to get there). While this standard definition in the literature, as just given, considers the reward uncertain and talks about expected (discounted) reward, I will not usually do that in this chapter for simplicity: I assume the world is largely deterministic.

<sup>5</sup> A multitude of algorithms for learning the state-values exist; see Sutton and Barto (2018) for a comprehensive treatment. To begin with, it is possible to use (stochastic) gradient descent on the objective function given by the total expected future reward. However, it is more common to use algorithms that proceed by a recursion where the value of a state is defined based on the values of the states to which the agent can go from that state. This recursion is based on the theory of dynamic programming, and in particular what is called the Bellman equation. The recursion basically ends at the goal state in the sense that the value of the goal state is given by the fact alone that it is a goal state, and the values of the other states are derived from that by the Bellman equation. Consider a simple world with three states, A, B, and C, where C is the goal state. You can move from A to B and from B to C. The first part of the recursion says that the value of state A must be the value of state B minus a small quantity. That is because from state A you could go to state B in a single step, and the small quantity expresses the fact that you need one step; it is a consequence of discounting. Likewise, the value of state B must be a bit less than the value of C. Now the value of C is fixed (to some numerical value which is irrelevant) by the fact that it is the goal, and needs no recursion or computation. So, once the agent encounters the state C even once, it knows the value of C. Based on that knowledge and its model of the world, it can start recursive computations, by applying the ideas above (value of B equal to value of C minus a small constant, value of A likewise) to recursively compute the values of B and A. If we fix the value of the goal to 1, the state-values could be 0.8, 0.9, and 1 for A, B, and C respectively. Note that in this example, we assumed the agent follows the optimal policy, i.e., it always takes the smartest possible actions; thus, the state-values computed are the state-values of the optimal policy. You could also compute the state-values for a very dumb policy (say, always taking random actions), and they would be lower because by taking less smart actions, the agent would get less reward.

<sup>6</sup>Again, strictly speaking, the action-values depend on the policy of the agent, and sometimes the term is used to mean the action-

really easy and extremely fast: Just compare the values of different actions and choose the one which gives the maximum. In fact, all the relevant information about the effects of the agent's actions are implicitly included in the action-value function. The agent still has to learn the action-values, but that is not really more difficult than learning the state-values.<sup>7</sup>

At the end of the 19th century, Edward Thorndike put cats in a box where they would need to press a lever to get out of the box and receive some fish to eat. He observed that in successive trials, the cats were pressing the lever more and more often. Such learning is called *instrumental conditioning* (as opposed to classical conditioning as in the famous Pavlov's dogs, to be considered below). This shows how learning to choose actions is possible by simply associating what we call a state in AI theory (here, being in the box) with an action.<sup>8</sup>

So, using reinforcement learning, an AI or an animal can actually learn to act without doing any real planning and having almost no model about the world. If it learns the action-value function, it only needs to look at the single actions immediately available, and then take the action which has the largest action-value—at the state where it happens to find itself. Since the action is here triggered immediately without any deliberation, like a *habit* or a knee-jerk reaction, the resulting behaviour is often called habit-based, or reactive.<sup>9</sup>

Reinforcement learning has recently become popular as a model of human behaviour in neuroscience, where humans may not be considered too different from experimental animals such as cats or rats. Current thinking is that the same reinforcement learning algorithms can be used to model at least one part of the action selection system in most animals, including humans. Nevertheless, there is little doubt that some animals, probably most mammals, engage in planning as well.<sup>10</sup>

In fact, reinforcement learning using value functions is not a magic trick that will obliterate the complexity of the action selection: It simply shifts the computational burden from search in the tree to learning a value function. Sometimes, this is a good idea, but not always. We will discuss the pros and cons of reinforcement learning vs. planning in Chapter 7. Let me just mention here the main disadvantages of habit-based behaviour:

---

values for the optimal policy. The terminology is further confounded by the fact that sometimes action-values can refer to the current estimates of the agent for those action-value instead of their true values (the same holds for state-values as well). Note that action-values are often called “Q-values”.

<sup>7</sup>For a comprehensive treatment of different algorithms, see (Sutton and Barto, 2018). Most algorithms for both action-value and state-value learning are based on very similar recursive computations based on the Bellman equation just described in footnote 5.

<sup>8</sup>This can be seen as an example of using something like action-values without a sophisticated world model. On the level of neurobiology, such reactive behaviour can also be explained by a special form of Hebbian learning, which implements something similar to the abstract theory of reinforcement learning we have just seen. In such learning, the association weight between the state (being in the box) and action (pressing the lever) increases every time both the state and the action are active, *and* a reward (fish) is delivered. Ordinary Hebbian learning would only be able to learn, in an unsupervised manner, the connection between the state and the action if the same action is frequently taken in a particular state. It would be useless in itself for selecting the *best* action since it does not take the reward gained by the actions into account. So, an extension of Hebbian learning to such “three-factor” learning, modulated by reward, is necessary. This may not be exactly what happens in the brain, but it is probably a useful approximation nevertheless (Nevin, 1999). Such three-factor (or modulated) Hebbian learning rules have a long history, see e.g. the discussions by Legenstein et al. (2010); Gerstner et al. (2018). These learning rules can also be extended to choosing action sequences in a dynamic environment: Basically, instead of the reward itself, the Hebbian rule might be modulated by reward prediction error considered next in the main text.

<sup>9</sup>Some authors use the term “model-free reinforcement learning” to clearly distinguish this from anything using planning. Planning uses a model of the world, thus it would be called “model-based”. Model-based reinforcement learning then refers to a set of algorithms for solving the planning problem, with the possible modification that instead of reaching a single goal, the plan may still attempt to maximize the sum of rewards.

<sup>10</sup>For a review on applications of reinforcement learning to modelling animal and human behaviour, see Niv (2009). On planning in animals, see Redshaw and Bulley (2018); Corballis (2019).

such learning often needs a lot of time and data, and leads to inflexible behaviour. This is quite in line with the common-sense idea we have about habits.

## Frustration as reward loss and prediction error

We have thus divided action selection into planning and habits, where habits refer to more automated action selection mechanisms.<sup>11</sup> Now we consider how to define frustration in the case of habits, where there are no goals but rather rewards obtained here and there, so we cannot talk about frustration in the sense of not reaching the goal.

What defines frustration in this case is an error signal called reward loss<sup>12</sup> which we already saw briefly in Chapter 2. It is computed by the following simple formula:

$$\text{reward loss} = \text{expected reward} - \text{obtained reward}$$

which is set to zero in case the difference is negative. That is, a reward loss is incurred when an agent expects to get some reward but actually gets less reward than expected. Maybe a cleaning robot expected to find a lot of dust in a room, but in fact there was much less. If it happens that the obtained reward is actually larger than expected, there is obviously no reward loss, so the reward loss is defined as zero if the difference in the expression above is negative. Reward loss can also occur if the expected reward is negative, and the obtained reward is negative while even larger (in absolute value): the agent did expect something bad to happen, but it turned out to be even worse.<sup>13</sup>

Expectation of reward here refers to the mathematical expectation as defined in probability theory. It is obtained by weighting the possible values by their probabilities: if the probability of obtaining a reward is 50% and the reward is 10 pieces of chocolate, the expected reward is 5 pieces of chocolate.<sup>14</sup> Expectations of the future are often called predictions, which are in fact ubiquitous in the brain: it is likely that the brain makes a prediction of almost any important quantity in the environment.<sup>15</sup>

<sup>11</sup>(Dolan and Dayan, 2013)

<sup>12</sup>(Papini et al., 2015; Mee et al., 2006)

<sup>13</sup>The effort made in trying to obtain the reward may also need to be taken into account in computing frustration. While it might seem natural to simply subtract the effort spent from the reward, considering it as a “cost”, sometimes more effort leads, paradoxically, to higher perceived reward (Inzlicht et al., 2018).

<sup>14</sup>The exact definition of expectation as used in reward loss is not very clear in my view, and an important problem for future research. Not much attention has been paid on it, partly because in typical experiments, it seems obvious what the expectation should be, and there is little planning involved. In a prototypical experiment, an animal (or a human) is given the same (positive) reward several times for a given, simple behaviour, and then suddenly it is given less reward (this is called “successive negative contrast”). In such a case, the future expectation of the reward is simply assumed to be equal to the past reward. With longer plans in more complex environments, the definition will be less obvious. Clearly, there is a strong connection to the concept of a prediction, as discussed next in the main text as well as footnote 15 below. Furthermore, an alternative definition of reward loss might be developed using counterfactual contrast (Roese, 1997), formalized as counterfactual regret by Zinkevich et al. (2008), where the obtained reward is compared with what *might* have been obtained, if better actions had been chosen. If the agents form some kind of a society, even more options exist for defining the expectations. The agent might use information on what other agents get, and expect to obtain the same reward as others do. Such a “social” expectation might simply be based on probabilistic inference: If the other agents are similar to the agent in question, it is logical to expect that the agent in question will be able to obtain the same amount of rewards (Rutledge et al., 2016). Yet another, very different, form of expectation might be produced in a situation where the agent assumes a *moral* right to obtain something, assuming the existence of some ethical norms in the agents’ society (Dignum et al., 2000).

<sup>15</sup>(Clark, 2013). From the viewpoint of mathematical theory, it might actually be more appropriate to talk about *predicted reward*

In contrast to our basic definition of frustration in Chapter 3, which works only on the level of plans, such a reward loss can be computed after every single action and at every single time point. This definition of reward loss is, in fact, quite flexible since the time interval in which the reward is computed can be specified to be anything from seconds to days. Therefore, it provides a general framework encompassing both planning and habit-based action. Reward loss coming from planning and reward loss coming from single actions are similar except that they work on very different time scales. We shall consider this point in more detail at the end of Chapter 7.<sup>16</sup>

Reward loss, in its turn, is related to what is called the reward prediction error (RPE), a most fundamental quantity in machine learning theory. RPE means any error made in the prediction of the reward. This definition is very general because the expected reward can be greater or less than the obtained one, and thus RPE can be positive or negative. If the obtained reward was larger than expected, that is the opposite of reward loss and suffering, and related to pleasure.<sup>17</sup>

As the very expression “reward prediction *error*” indicates, the theory of RPE also shows how suffering is related to learning by minimization of errors, which is a fundamental approach in machine learning. If the agent can predict the rewards obtained by different actions in different states, it will be able to act so as to maximize the obtained rewards (at least if it has a good model of other aspects of the world as well). To learn and improve such predictions, it is necessary to compute the errors incurred by the current system to predict the reward. This is how minimizing RPE is related to maximization of rewards. In fact, it is possible to devise reinforcement learning algorithms that work simply by minimizing reward prediction error.<sup>18</sup>

The exact mathematical definition of RPE is quite involved and relegated to a footnote.<sup>19</sup> Let me just point

---

instead of expected reward in the definition of reward loss. While these are often seen as the same thing—prediction being an expectation of a future quantity—the concepts are not equivalent. In particular, in machine learning theory, a prediction can be considered more general than expectation: a sophisticated prediction will also include an estimate of the uncertainty involved in the prediction, in addition to the mathematical expectation. This is relevant here because it seems that the certainty of the prediction affects the level of frustration. I would claim that if you are completely certain that you will get chocolate (say, 5 pieces), but then it turns out you don’t, the frustration will be greater than in the case where there is only some chance of getting any (like the example in the main text, 10 pieces with 50% probability). Crucially, in this example the expected amount of chocolate, in the sense of the mathematical expectation, is the same in the two cases, and only the uncertainty changes. Such an effect of uncertainty should be taken into account in the definition of reward loss. I will not do that completely rigorously in this book because such theory seems to be lacking at the moment; however, closely related developments on uncertainty will be found in Chapter 10 and footnote 19 in Chapter 14 (see also the text preceding that footnote).

<sup>16</sup>Another difference is that while earlier (Chapters 2 and 3) we defined frustration as “not getting what one *wants*”, in line with the quotes from ancient philosophers, here reward loss is defined as “not getting what one *expects*”. These are not exactly the same thing and are sometimes quite different; this connection will also be discussed in Chapter 7, page 86.

<sup>17</sup>I refrain from trying to rigorously define pleasure in this framework, but obviously an RPE where the obtained reward is greater than expected is a good candidate. In fact, in experiments with participants playing a gambling game, reward prediction error was a strong predictor of the participants’ well-being (Rutledge et al., 2014), presumably including both pleasure and suffering symmetrically; however, the average level of reward had a strong effect as well. Alternatively, Carver (2003) has proposed that the function of the pleasure system is to signal that the current task has been accomplished, and the system can direct its resources to other tasks. The neurobiology of pleasure and pain is reviewed by Leknes and Tracey (2008).

<sup>18</sup>This can be done using a special form of RPE called temporal difference (TD) error. See Sutton and Barto (2018, p. 268) who call it Bellman error, or related developments by Bhatnagar et al. (2009).

<sup>19</sup>RPE can actually be defined in different ways. In neuroscience literature, the definition may not be very different from reward loss. In the reinforcement learning theory, a more sophisticated definition is usually used, using what is specifically called the temporal difference (TD) error, which we explain here. We consider the case where a single action is taken at each time step (as opposed to considering longer plans), and no discounting is used. For each time step, RPE is then defined as  $RPE = \text{reward} - (V_{\text{before}} - V_{\text{after}})$  where  $V$  is the state-value function (for the policy being followed, not necessarily the optimal one), in the state before the action was

out that RPE is a more general concept than reward loss also in the sense that it enables an error signal even when the agent is far away from any actual or expected rewards, but it receives “bad news” about future reward. This is in contrast to reward loss which does not make any sense except when a reward is actually expected to be obtained right now (the exact meaning of “now” depends on the time scale). In particular, if the expectation of total *future* reward decreases, RPE signals frustration, while the reward loss has no meaning if no reward was expected to be obtained at this moment. Suppose a cleaning robot is on its way to a room where there is a lot of dust (yummy!), and thus its expected (predicted) reward is high, but this reward will in any case not be obtained for quite a while. Then, it finds that the door to the room is locked and it cannot enter; that is bad news it didn’t expect. Thus the robot finds itself in a new state that has a much lower expected total future reward since the dust in that room cannot be reached. It is this difference between the earlier prediction and the new prediction that creates RPE and suffering. This is not an ordinary reward loss because no actual reward was expected to appear at this time point anyway: the robot has not yet even entered the room, and the dust is still far away. However, RPE can create suffering merely based on predictions: if information arrives that makes the agent reduce its prediction of future reward, frustration is created. This is intuitively appealing since a lot of our frustration is actually about such negative news and the lowering of expectations they create. Suppose I’m planning to attend an event that I expect to enjoy, and then, well in advance, I hear the event has been cancelled. I will suffer, although I didn’t expect to obtain anything enjoyable yet, and I may not have taken any action either; it was all just predictions in my head.<sup>20</sup>

taken or after the action taken, respectively (which could also be denoted by time indices  $t-1$  and  $t$ ). The reward is the reward obtained for this particular action, or in other words, at this particular time step for which we are computing the RPE. Note that the sign is flipped compared to the definition of reward loss, but this is just a technical convention with no deeper meaning.

The connection to reward loss can be seen by understanding that in the state-value formalism,  $V_{\text{before}} - V_{\text{after}}$  can be interpreted as expected reward. The reason is that by the definition of the state-value function, the state-value function gives the total reward expected when starting from each of the states, so you would expect a reward equal to  $V_{\text{before}} - V_{\text{after}}$  to appear. Otherwise the two state-values would be inconsistent; the total expected reward starting from the state “before” must be equal to the total expected reward starting from state “before” plus the expected reward obtained in the transition. So, the agent can expect that reward =  $V_{\text{before}} - V_{\text{after}}$ , and if that actually holds, RPE would be zero. If you get less, there is a reward loss, which is here expressed as a negative RPE.

Such an RPE signal is the foundation for reinforcement learning. It is more general than reward loss since it considers the whole future of rewards via the state-values, as explained next in the main text. There are also some small differences: RPE has a different sign, corresponding to the negative of reward loss, and our definition of reward loss considers only the case where it is positive (or RPE is negative) since this is the part corresponding to suffering. Also, typically the discounting formalism of reinforcement learning is included in the definition, in which case  $V_{\text{after}}$  would be multiplied by a discounting factor throughout; omitting the discount factor is possible if we consider a finite time horizon. See Sutton and Barto (2018, Ch.15) for more information.

<sup>20</sup>To see how this works mathematically, consider the definition of RPE (given in footnote 19 above) in the case where the obtained reward is zero. It is in fact generally agreed that rewards are sparse, often extremely sparse, so most of the time the RPE is simply the difference between the state-values in two states (before and after, or past and present), possibly discounted in the latter state. Recall that the state-value is nothing else than the predicted total future reward. Thus, recalling that the sign in this conventional definition of RPE is wrong for our purposes, RPE defines frustration as  $V_{\text{before}} - V_{\text{after}}$ , which is exactly the decrease in predicted total future reward, comparing the prediction in the previous time step and the present time step. Such a decrease is possible when the agent receives new information (in the basic formalism, it finds itself in a new state), and that information makes it update its prediction (it switches to the prediction given by the new state it finds itself in). To summarize, in the case where the prediction decreases in the absence of any reward obtained, the negative part of RPE is thus the decrease in the prediction of the total future reward, or more precisely: previously predicted total future reward minus currently predicted total future reward. This is how RPE can define frustration based on predictions alone, without any reward currently expected. One might think that reward loss could do the same if we simply change the time scale: in the robot example, if you take the expected and obtained reward for, say, one whole hour, that would arguably lead to a reward loss since the robot expected to get dust during that hour but didn’t get any. However, RPE makes its computations independently of any such time scales (it is in fact taking into account the whole future as it looks at the total expected future reward) and moreover,

## Expectations or predictions are crucial for frustration

Reward loss and RPE highlight the importance of expectations and predictions. Clearly, there must be some expectation or prediction in order for them to occur. If the cleaning robot were so primitive that it had no expectations or predictions at all, it might be just enjoying every single speck of dust it finds. Making it more intelligent so that it can predict the future thus deprives it of its “innocence”, and enables frustration to occur. Likewise, Cassell says that “to suffer, there must be a source of thoughts about possible futures”, even though his approach to suffering is quite different.<sup>21</sup>

The importance of predictions is well appreciated in neuroscience. It has been observed that in the brain, RPE is coded by certain neurons using a neurotransmitter called dopamine. More precisely, it is coded by quick changes in the level of dopamine (called “phasic dopamine signal”), typically originating in evolutionarily old areas such as the midbrain, which is literally in the very center of the brain.<sup>22</sup> In case the obtained reward is higher than expected, there is a temporary peak in the amount of dopamine in the signalling pathways, which is called by some a “dopamine surge”. That’s why many drugs of abuse target the dopamine pathways in the brain. For example, cocaine blocks the removal of dopamine in the synapse so that its signal is amplified.<sup>23</sup> Such drugs are fooling the reward-processing system in the brain, thus leading to a strong desire for such drugs, in addition to a pleasurable feeling. This has led some to think that dopamine is the neurotransmitter responsible for the feeling of pleasure itself. Such a viewpoint is probably incorrect, and the actual feeling of pleasure is mainly mediated by other transmitters, namely those in the opioid family, while dopamine is more related to “cold” action selection and learning.<sup>24</sup>

### *Classical conditioning*

To emphasize the importance of predictions in the brain, let’s consider an extremely famous kind of prediction learning in the animal realm: classical conditioning. Ivan Pavlov, doing physiological experiments on dogs around the year 1900, observed that the dogs began to salivate when they saw the staff person who was responsible for feeding them, even before receiving any food. Pavlov was intrigued and tried to see if the dogs would be able to associate any arbitrary stimuli to food. He succeeded in making the dogs associate food with many different kinds of stimuli, including the sound of a bell or a metronome, provided that these stimuli were

---

such long-term reward loss would not occur before than hour has passed, while RPE signals frustration the very moment the new information has arrived and has been processed. (As a minor point on terminology, it may be slightly misleading to talk about “reward prediction *error*”, since RPE is in this case rather a change in predictions due to new observations; a non-zero RPE does not necessarily imply that there was any error, but simply a change, an update of prediction based on new information. )

<sup>21</sup>(Cassell, 1989). Likewise, the importance of predictions and expectations in economic decision-making is emphasized by Köszegi and Rabin (2006) who propose that consumers compare expected utility given an action with a “reference-point” given by a probabilistic prediction of the future utility.

<sup>22</sup>(Schultz, 2016). In experiments with humans, such signalling might be measurable as the error-related negativity (ERN) seen in EEG measurements, as well as fMRI signals mainly in some parts of the anterior cingulate cortex where ERN seems to originate (Holroyd and Coles, 2002; Abler et al., 2005).

<sup>23</sup>(NIDA, 2020), but see also Nutt et al. (2015)

<sup>24</sup>(Berridge and Kringelbach, 2015; Leknes and Tracey, 2008). This dissociation may sound logically contradictory, but it is based on the distinction (in Berridge’s terminology) between the motivational “wanting” processes which more directly tell the organism what to do, and the affective “liking” processes which are related to the feeling of pleasure. Abler et al. (2005) also proposes that reward loss triggers different kinds of neural processes, some of which are more related wanting, action selection and reinforcement learning, and others more to liking and the feeling of pain or pleasure; they find that the localizations of those two processes in the brain are different.

consistently presented just before food was given.

What the animal is clearly doing here is predicting the future: After the bell, food is likely to arrive. Such predictions are ubiquitous in the brain; the brain is constantly trying to predict what happens, using many different systems. Predicting the results of any action you might take is important if you want to choose good actions, as we already saw in the case of instrumental conditioning above (page 51). Predicting where the rabbit will be a second or two later is necessary if you want to catch it. Note the crucial distinction between classical conditioning and instrumental conditioning: In classical conditioning, the agent does not yet learn to *choose* actions, but merely to predict future states, independently of any rewards.<sup>25</sup>

It would be natural to assume that such classical conditioning could be easily performed by Hebbian learning. It is just the kind of association of two stimuli—bell and food—that Hebbian learning seems to be good at. That is to some extent true, although this is a bit tricky; the most successful models actually use supervised learning, with the bell as input and the food as the output. Such learning, again, proceeds by minimizing prediction error.<sup>26</sup>

### *Does a low level of rewards produce suffering?*

Intuitively, however, it might seem that talking about frustration based on expectations and predictions is unnecessarily complicated. If the agent is in a state with low state-value (in its own estimation), would that not in itself imply suffering? Being in a state of low value means that the agent believes it will not obtain much net reward in the future, which sounds like a good reason for mental pain. Or, even more fundamentally, why not just say that lack of rewards, presumably during recent history, is suffering?

One fundamental problem with such an approach would be that it is not obvious how to define a suitable baseline or comparison: What level of state-value is actually low, and how small should recent reward actually be to create suffering? The reward loss or prediction error actually solves this problem by using the expected reward as the baseline. Thus, the obtained level of rewards is compared with the expected level, and if it is “low” in this particular sense, suffering occurs.<sup>27</sup>

<sup>25</sup>In Pavlov’s experiment, the dog learned to predict that food is coming, independently of its actions. It did salivate, which could be seen as an action, but the salivation was a (presumably innate) response to food that was not learned during this experiment.

<sup>26</sup>For a single conditioned (i.e., predictive) stimulus, Hebbian learning actually works fine, but the problem is that when there are several conditioned stimuli, Hebbian learning would create too many associations and in an unbalanced way. For example, we could have an experiment where both a bell and a green light predict food. Simple Hebbian learning would then associate both those stimuli with the food, and the association strengths would be computed independently of each other; since the association strengths are computed independently of each other, the predictions may interfere with each other and lead to bad prediction. This has been investigated in a famous twist to the basic classical conditioning experiment: after the main experiment, another experiment is made where both the bell and a newly introduced green light predict food. In such a case, the dog will not learn to associate the green light with the food because the connection from the bell is enough to predict the food, and there is no need to construct an association from the light to the food anymore. This is in contrast to what Hebbian learning is supposed to do. The brain apparently tries to be economical and constructs only those connections that are necessary for the prediction of the food. Therefore, the association strength of one conditioned stimulus will also depend on the associations of other stimuli. This is why most research assumes a supervised model, which typically learns several such association strengths in a balanced way, and thus explains the various experiments better than simple Hebbian learning. A basic supervised learning rule accomplishing this is the Rescorla-Wagner model; see e.g. review by Miller et al. (1995). It actually further models the dynamics of learning, as in the bell/light example above, where it is important that the bell is first associated to the food and the light only comes later; the association with the bell “blocks” the development of a new association with the light.

<sup>27</sup>The RPE formalism could also be interpreted as providing another baseline mechanism, by looking at the *change* of state-values. Going to a state which has a lower value than the current state, without obtaining any reward, does produce suffering according to



## Unexpected implications of state-value computation

In the rest of this chapter, I will consider some practical implications of the theory presented here. First, let us consider how the computation of state-values, as proposed in basic reinforcement learning theory, fundamentally changes the behaviour of human agents. Originally, of course, evolutionary forces demand that an action is pursued by a biological organism if it helps it reproduce and spread our genes, and an action is avoided if it hampers this effort. So, evolution “tells” us that kicking a stone is bad because it can cause damage to our foot, and the damage decreases our potential for reproduction—thus giving us negative reward for such an action. Having sex is very good, and rewarded by basic evolutionarily mechanisms, because then we are fulfilling our deepest evolutionary calling and spreading our genes.

The computation of state-values changes the situation: The organism will not only try to reach states directly giving reward—such as having sex—but also states that have higher state-values. This is a mechanism for looking forward in time: instead of immediate reward, the organism will try to maximize the total reward in the future, and just that is given by the state-value.

Seemingly valueless states are now valued by the agent since they *predict* that more actual reinforcement can be found sooner. Such states provide intermediate goals in the pursuit of the actual reward, similarly to heuristics in tree search. If you train a robot to get orange juice from the fridge, it must of course first go to the fridge, and open it. So, the state where the robot is standing next to the fridge acquires a positive state-value and we could even say that the robot “likes” being next to a fridge, even more so if it is open.

The situation is even more complex due to the existence of human civilization and society. Culture plays an important role in determining the state-value function, and it is often difficult to separate the influences of biology and culture. In neuroscience, this is called the “nature vs. nurture” question. There can be extremely complex chains of value computation which transform the original evolutionary goals to behaviour based on intermediate goals. For example, humans have evolved to strive for high social status. From an evolutionary perspective, this is because it helps humans get more sexual partners and increases the number and the survival probability of their offspring. This then implies that we want to increase our status: for example, winning a gold medal in the Olympics is a good behavioural goal. Clearly, a gold medal in Olympics has no evolutionary value in itself: it does not satisfy your hunger, thirst, or sexual appetite in itself. It is just an arbitrary piece of metal. There is no logical connection between such a piece of metal and sex. It is only due to a complex interplay of value function calculation *and* cultural meanings that the original evolutionary reward of sex has been subtly transformed into a goal such as excelling in sports—or science, or politics.

Such slightly weird desires are another manifestation of the phenomenon discussed earlier: emergence of unexpected phenomena due to the interaction between the learning agent and a complex environment. If we program sufficiently complex AI, the same thing is likely to happen as with human evolution. The AI will pursue goals that were not intended by the programmers, but which still happen to produce a high state-value.<sup>28</sup>

---

the definition of RPE above, as explained in footnotes 19 and 20 in this chapter. In this sense, RPE uses the current state-value as the baseline defining what is “low”. See also footnote 14 on different possibilities of defining the baseline as “expectation”.

<sup>28</sup>Human society, with its complex interactions between individuals and groups of individuals, creates a particularly complex environment. When AI’s interact not only with humans but with each other, that will add yet another layer of complexity, and quite unexpected things might happen.

## Evolutionary rewards as obsessions

Now, if we admit that our desires are based on evolution, even if quite indirectly, is that a good thing or a bad thing? Should we just follow our desires, or think twice, or even try to follow some completely different goals? There are actually people who try to justify certain kinds of behaviour by saying they are evolutionarily conditioned, i.e. “evolution made me do it”. In popular science magazines and web sites, such logic is not very uncommon. Fortunately, it is rejected by many as an example of sloppy thinking.<sup>29</sup> In the following, I argue the very opposite: following evolutionary desires is often a bad idea and even morally wrong.

In fact, even the evolutionarily conditioned *rewards* themselves can go wrong, sometimes quite catastrophically. One reason is that evolutionarily, we may be adapted to the environment where our evolutionary ancestors lived, often assumed to be the “African savannah”. However, the modern world is different and, therefore, our evolutionary programming may not be very suitable.<sup>30</sup> With humans, a well-known example is the addictive quality of sugary food. The sweet taste of sugar must have signalled the high nutritious quality of food in the environment where our ancestors lived.<sup>31</sup> But these days it tends to signal added refined sugar which is bad for your health; evolutionarily speaking, sweet taste should rather be punishing in the modern context, not rewarding.<sup>32</sup> Yet, the state of having a sweet taste in your mouth is rewarding, and humans tend to try to reach such a “sweet” states.<sup>33</sup>

What is even more serious is that evolution makes us want particularly questionable goals, especially from a societal viewpoint. Evolution is fundamentally based on selfish, merciless competition between different organisms (or strictly speaking, between their genes). Many behavioural tendencies evolution has imposed on us should be seen as instruments for such egoistic competition. Evolution is all about maximally spreading our genes. It makes us hoard finite resources such as food to ourselves in order to spread our genes. It makes us violent; it even makes us go to war, again for the sole purpose of spreading genes. This is in stark contrast to most ethical systems in the world which see such selfishness as evil, and recommend quite opposite courses of action.<sup>34</sup>

Even more fundamentally, the rewards defined by evolution never had the goal of making us happy in any

<sup>29</sup>It is a case of what G.E. Moore called the naturalistic fallacy. Hume already pointed out that you cannot infer what *ought to be* from what *is*. In other words, if evolution makes people behave in a certain way, it does not in any way *morally* justify the claim that this way of behaving is good or acceptable.

<sup>30</sup>(Sapolsky, 2004; Wright, 2017)

<sup>31</sup>Another striking example in the case of humans is pornography, where watching sexually desirable models on a computer screen is felt to be somehow rewarding by its human consumer, and leads to desire towards such pictures. Simply seeing sexually attractive people naked should indeed have a very high state-value, since that is likely to happen only when copulation is near—at least in our evolutionary past. But in the modern world, this behaviour is quite dysfunctional in the sense that there is almost no chance that the consumer would be actually able to mate with those models.

<sup>32</sup>As a kind of mirror image of such maladaptive evolutionary desires, there is the phenomenon of chronic (persistent) pain. Raffaelli and Arnaudo (2017) review research on how chronic pain “entails a pathologic reorganization of the neural system” so that it “loses its biologic damage signaling function” and “becomes a destructive force”, eventually a disease in its own right.

<sup>33</sup>In fact, it is often difficult to define what is the actual reward and what is differences in state-values. I’m here assuming that the sweet taste is a reward in itself, and not a question of a high state-value (i.e. predicted future reward), but this can be disputed. It is less controversial that an Olympic medal does not produce a reward in itself, but even this is not so clear. To solve this problem, Singh et al. (2009) propose that the rewards should evolve so that they are correct in most environments, while state-values are then learned during an individual’s lifetime for the particular environment where the individual is living.

<sup>34</sup>Admittedly, the connection between ethics and evolution is complex, and evolution seems to have conditioned some kind of altruism in us as well (Wright, 1994; Nowak et al., 2010).

meaningful sense. They are a force that drives us to do exactly those things which are good for our evolutionary fitness. Even if you come to the conclusion that the evolutionary reward system makes you suffer, it cannot be switched off or modified. You cannot decide to be rewarded by something you consider more meaningful and good for society.

I think what evolution offers us is something I would call *evolutionary obsessions*.<sup>35</sup> That is, the evolutionary rewards, together with the learned state-values, make us desire, even crave for many things which we would actually prefer *not* to desire if we could rationally decide what we desire. If you could just consciously, rationally, “switch off” your desire for, say, sugary food—would you not do that? Chapter 14 explains how Buddhist and Stoic thinking are based on the rather extreme tenet that switching off *all* desires would actually be very good for you. Whether one agrees with that extreme viewpoint or not, surely, most people have certain desires that they would rather not have. I call them obsessions because they are automatically created, they often override any conscious deliberation, and they may even feel unwanted and intrusive. (We will look at the computational mechanisms for this in Chapters 7 and 8.)

## Reward maximization is insatiable

Finally, let me mention another dark side to this reinforcement learning theory. One crucial property of the algorithms based on reward prediction error is that they drive the system to get more and more reward, and there is never any long-term satisfaction. This is because any prediction of the future is learned by the agent, and *constantly updated by learning*. Thus, in the reward loss, the level of expected reward is updated based on what the agent has obtained recently.

Suppose that an agent gets an exceptional amount of rewards for a while, maybe because a cleaning robot finds itself in a building with lots of nice dust to clean, and it is rewarded for every speck of dust it sucks away. Now, the agent’s prediction system is updated so that an equally large amount of rewards is predicted in the future as well. An environment that produced an unexpectedly large amount of reward for a while becomes the new baseline. That level of reward is not unexpected anymore and, therefore, does not produce any particular “pleasure” anymore either.<sup>36</sup> What’s worse is that when things get back to normal, the agent will get less rewards than what it has now learned to expect, since the prediction was updated to reflect the particularly nice environment that lasted for a while. Therefore, the agent suffers enormously when it has to go back to a normal room with a modest amount of dirt.

Similar computations take place in our brain, since our brain also computes the reward prediction error and updates its expected level of reward. No wonder that Wolfram Schultz, one of the leading neuroscientists on dopamine, calls the dopamine neurons “little devils”.<sup>37</sup> In fact, this is a logical consequence of the guiding principle of AI agent design: the agent should maximize obtained reward. The reward prediction system has

<sup>35</sup>I am here using the term “obsession” in a loose sense, not using its strict psychiatric definition. For reference, in the current ICD-11 proposal, obsessions are defined as follows. “Obsessions are repetitive and persistent thoughts (e.g., of contamination), images (e.g., of violent scenes), or impulses/urges (e.g., to stab someone) that are experienced as intrusive, unwanted, and are commonly associated with anxiety. The individual attempts to ignore or suppress obsessions or to neutralize them by performing compulsions. — Compulsions (or rituals) are repetitive behaviours (e.g., washing, checking) or mental acts (e.g., repeating words silently) that the individual feels driven to perform in response to an obsession, according to rigid rules, or to achieve a sense of ‘completeness’.” (Stein et al., 2016)

<sup>36</sup>This is related to the phenomenon of the “hedonic treadmill” (Lyubomirsky, 2010).

<sup>37</sup>(Schultz, 2016)

no other goal than helping in maximization of rewards. If you program an agent to maximize reward, then by definition, nothing can possibly be enough; the system will be *insatiable*. The agent will relentlessly try to get more and more reward, and it is precisely the frustration signal that will force the agent to try harder and harder.<sup>38</sup>

A merciful programmer might program some stopping criterion to limit the greed of the agent: Once you have obtained X units of reward, you can stop. Unfortunately, evolution knows no mercy, and humans don't seem to have any such stopping criterion programmed in them. We need more money, more power, more sex (and better sex), and better food (and more food). If we follow our evolutionary “obsessions”, as I called them, nothing is enough.

Suppose you program a robot called Pat to clean a building. You would like the building to be superclean, and the building is quite large with dozens of rooms. So, you would be very tempted to program Pat so that it will spend all its time cleaning the building. You probably want to program a couple of other functions in Pat as well, such as a routine for charging its batteries, some basic maintenance procedures, as well as safety systems to prevent it from hurting people or breaking things. But you would probably program Pat to spend all the rest of the time in tirelessly cleaning the rooms, with no breaks in between. This is what most programmers would do. Here, you have implemented a kind of a “cleaning drive” which is without mercy. Pat will spend all its time and energy just making the rooms spotlessly clean. This may seem completely natural, given that it is “just” a robot.

Now, suppose your colleague, responsible for the visual design of the robot, decides to make Pat look really cute, giving it the shape of a little kitten. It even says “Meow” using its loudspeakers. Many people may suddenly start feeling sympathy for this poor little kitten. “Does it really have to be working all the time? Can't it ever play, or take a rest?” they would ask. What would you reply?

---

<sup>38</sup>Lambie and Haugen (2019) consider insatiability as an important component of greed. On the other hand, it is true that some purely biological needs are satiable to some extent—for example, hunger is reduced by eating, even if momentarily—but classical reinforcement learning theory is lacking much consideration of such metabolic states (Keramati and Gutkin, 2014).

## Chapter 6

# Suffering due to self-needs

In addition to frustration, we identified another cause of suffering in Chapter 2: loosely speaking, suffering related to self. Self is a concept with a bewildering array of meanings. Psychology, philosophy, and neuroscience offer a multitude of definitions, and I can make no claim to treat the concept comprehensively.

I focus here on two meanings of “self” directly related to suffering. First, self as the target of evaluation of some kind of long-term success of the agent. The human brain, in particular, has a system that constantly evaluates the agent, checking whether the goals set were reached or rewards obtained, and seeks to improve its general performance. Second, we have self as the target of self-preservation, or survival instinct: all animals have behavioural tendencies to avoid death or organic damage. (A third meaning of self, related to control, will be treated in Chapter 11, and the concept of self-awareness, in Chapter 12.)

Such self-evaluation and self-preservation are computational mechanisms which are constantly operating in animals, and it is easy to justify their computational utility for any intelligent agent. Although at first sight, these aspects of self may seem to provide a mechanism for suffering which is completely different from frustration, I will also show how they are related to frustration of internal, higher-level goals and rewards. As the title of this chapter indicates, these aspects of self can thus be seen as needs, or desires, and they can be frustrated.

### Self as long-term performance evaluation

Let us start with self as something whose performance is being constantly evaluated at different levels. As we saw earlier, in reinforcement learning, every single action is always evaluated to improve future actions. The reward prediction error is computed even in the simplest algorithms. If the reward is incorrectly predicted, the error is used by the learning algorithm to improve the prediction—if the prediction was too high, set it lower in the future, for example. While such computations are crucial for learning to act optimally, the errors also trigger the suffering signal according to the theory of the preceding chapter.

However, the situation is complicated by the fact that the learning algorithms themselves contain many parameters describing how the algorithm itself works. One fundamental parameter is how quickly the system should learn: If it learns too quickly, the new information will tend to override the old one, thus leading to forgetting. Below, we will see another parameter which is how much of the time the agent should spend on relatively random exploration of the environment. There are many such parameters in a sophisticated learning system.

Therefore, sophisticated AI should be able to adjust such internal parameters by itself. This is called *learning to learn*.<sup>1</sup> Such learning to learn requires constant monitoring of the performance of the basic learning algorithms. If the current internal-parameter settings do not lead to good learning, adjustments have to be made. This requires an internal signalling system, not unlike the suffering signal, but typically working on a longer time scale, since it takes a long time to see if a learning system learns well.<sup>2</sup>

### *Self-esteem and depression*

In humans, mood is a signalling system working on a longer time scale. Mood is defined as an emotional state which is more long-lasting than single emotional episodes (such as being angry or feeling afraid). A low mood may take days, if not weeks or months, to change. A psychological concept which works on an even longer time scale is self-esteem: an overall view of the self as worthy or unworthy.<sup>3</sup>

Depression may in fact be an extreme case of the performance signalling made by the self-evaluation system. One theory proposes that depression occurs when goals are not reached, and moreover, constant attempts to improve performance fail.<sup>4</sup> That is, the agent has to admit that whatever it tries, nothing works. In such a case, there is still one last strategy that may help: wait and see. The environment may eventually change by itself, even if you do nothing. Perhaps, after a while, with some luck, the circumstances will be more favourable. Such a “wait and do nothing” program may explain some depressive symptoms, such as passivity and lack of interest in any activities.<sup>5</sup>

It would clearly make sense to program such a “depressive” mechanism in an AI. If the current algorithms are simply not working at all, it would be better for the agent to just wait and see if the world changes for the better. Such waiting will save energy, and perhaps will also enable the AI to perform some further computations to improve its performance in the meantime.

### *Self-destructing systems*

What if an AI comes to the conclusion that it is not able to fulfill its task at all? Perhaps something went very much wrong in the design of the learning algorithm, or the task is completely impossible, and the circumstances do not seem to change for the better. The most extreme solution would then be for the AI to “destroy” itself.

Suppose you launch many AI agents, or programs, that work more or less independently inside some computing system. If one of the agents is not achieving anything, it would be natural that you terminate its execution. This would free up computational resources for other agents—assuming all the agents are running on the same shared processors—and other agents might be more successful.

---

<sup>1</sup> (Thrun and Pratt, 2012)

<sup>2</sup> In humans and other social species, a related aspect of “self” is how oneself is seen by others (Sebastian et al., 2008; Heatherton, 2011). This is clearly related to an evaluation of the self, but in this case, performed by other agents.

<sup>3</sup> (Heatherton et al., 2003)

<sup>4</sup> (Thierry et al., 1984; Nesse, 2000)

<sup>5</sup> A very similar account has been proposed for the simple emotion of sadness by Oatley and Johnson-Laird (1987). The difference is mainly in the time scales involved, since depression is by definition much more long-term than sadness. Sadness in its turn could be seen as a frustration or disappointment signal which is particularly strong and relatively long-lasting, but the terminology here is not very well-defined. A related computational account of depression focusing on the concept of learned helplessness is given by Huys and Dayan (2009).

To make this possible, there has to be a system for evaluating each AI agent's performance as a whole. Importantly, the evaluation does not have to be done by an external mechanism; it could be part of the agent itself, which could then decide to self-destruct. There is nothing paradoxical or impossible in such a self-destruction system. It can be explicitly programmed in the agent by a human programmer—while it may indeed be quite impossible for the agent itself to learn such self-destruction behaviour.

It is possible that in some cases, even biological organisms may engage in such self-destruction sequences. Such an idea is quite speculative because it is not obvious why evolution would favour such behaviour. It is clearly possible that the designer of an AI system can explicitly create the self-evaluation and destruction systems, but in biological evolution, there is no such explicit designer. It may actually sound completely non-sensical to think that evolution could lead to self-destruction mechanisms, since an organism which destroys itself cannot spread its genes anymore.

However, evolution is a bit more complicated than just the survival of the fittest individual. It is widely appreciated that in evolutionary arguments, we should take into account not only the survival and reproduction of an individual, but also the survival and reproduction of the closest relatives. This leads to the concept of “inclusive fitness”, where the fitness of an individual takes into account the fitnesses of the relatives weighted by the proportion that they share genes. Close relatives of an individual spread partly the same genes anyway, so their survival is evolutionarily useful for that individual.

According to one suggestion, if a person is seriously ill, and finds himself a great burden to his relatives, it might actually be evolutionarily advantageous for that person to commit suicide. If this helps the relatives with whom he shares a large proportion of genes, the suicide might actually help in spreading those genes, thus increasing the inclusive fitness.<sup>6</sup>

Thus, self-destruction programs may be useful not only to maximize the utility of AI agents, but also from an evolutionary perspective. This may sound abhorrent from a moral perspective, but that is often the case with evolution which has no reason to be nice or good from a human perspective—as already argued in the preceding chapter, where I compared evolutionary desires to obsessions.

## Self as self-preservation and survival

Another rather obvious reason why some kind of concept of self should be programmed in an AI is that the AI may need to protect itself against anything that might destroy it. A robot must take care not to be run over by a car: This is the concept of self-preservation. There is no doubt we can, and probably want to, program some kind of self-preservation mechanism in an AI agent.

Even the simplest biological organisms have behavioural programs that are activated when their existence is threatened; we talk about self-preservation, or survival instinct. We already encountered related ideas in considering definitions of pain and suffering. The widely-used IASP definition related pain to “tissue damage”

---

<sup>6</sup>(de Catanzaro, 1991). However, see (Nowak et al., 2010) for a criticism of the centrality of kinship in the inclusive fitness theory. Related work on suicide and self, but without the evolutionary interpretation, is by Baumeister (1990). Taking the logic of inclusive fitness even further, one may be tempted to think of natural selection working on the level of groups of organisms (families, tribes, herds, etc.), so that it is the fittest group, not organism, that survives the selection. However, any theories based on such “group selection” are controversial, and it is not clear if it actually happens in nature. Some mathematical theories propose that natural selection on the level of individuals leads to emergence of phenomena which look just like the selection happened on the level of groups. In fact, according to the mathematical model by Hadany et al. (2006), something like self-destruction could actually emerge from purely individual-level selection, in cases when an individual organism finds the current environment particularly adverse.

(page 15), while Cassell's definition of suffering talked about the "intactness of the person" (page 17). However, what we are talking about here is threats to the very existence of the agent, not just damage.

While it seems relatively straightforward to program self-preservation behaviours in an AI, an open question is whether an AI can somehow develop a survival instinct by itself. In other words, can self-preservation emerge without being explicitly programmed; can the agent learn to perform certain actions for the main purpose of avoiding its own destruction? This is one of the deepest questions in AI, extremely relevant from the viewpoint of developing safe AI systems, and the subject of intense debate.<sup>7</sup> We have seen earlier that learning in AI can have various side-effects and unintended consequences; this would be one of the most extreme ones.

On the one hand, there are those who point out that biological organisms have developed their survival instinct via evolutionary mechanisms. They have been subject to natural selection, which has ruthlessly weened out those organisms which do not fight for their survival. In contrast—this line of argumentation goes—AI is not subject to natural selection; it has no evolutionary pressures. So, it will not learn a survival instinct, unless perhaps we explicitly decide to program it to learn one.

Other experts disagree and point out that some kind of survival instinct may be automatically created as an unintended side-effect of creating sufficiently intelligent machines. If a robot is given any mundane task, say fetching a bottle of milk from a near-by shop, a super-intelligent robot would understand that in order to perform that task, it has to stay alive. If the robot were damaged or destroyed in a collision with a car, for example, its task cannot be performed. Thus, the robot might decide to destroy the car somehow (let's assume the robot is really big) to get the milk safely delivered. If everybody in the car gets killed, that is irrelevant, if the programmer didn't tell the robot to avoid human casualties. The idea here is that there is no need to explicitly program a survival instinct, or any reward related to that: the general goal of maximizing future rewards will direct the robot's behaviour towards avoiding destruction. In fact, this line of thinking means that almost any sufficiently intelligent AI will *by logical necessity* strive to survive. If it is intelligent enough, it will understand what death is, and how death makes it impossible to obtain any further rewards or accomplish goals. This is the opposite of what has happened in biological evolution, where even the very simplest organisms have a survival instinct, and sophisticated intelligence develops later. In AI, intelligence is programmed first, and later, possibly by chance, the AI might obtain a tendency for self-preservation behaviour and related information processing, which might then be called a survival "instinct".

Clearly, these two views are based on very different assumptions about the AI. The argument where the robot understands that a car on a crash course has to be destroyed assumes a very, very intelligent robot. The robot must have a sophisticated model of the world, infer that it risks being overrun by the car, and understand that being overrun by the car will prevent it from delivering the milk. Most current robots would be nowhere near the intelligence required—but we don't know if they will be in the future. We are even further away from an AI which could intellectually infer, on an abstract level, that there is such a thing as death, and that various measures should be taken to avoid it.

Nevertheless, if an AI is learning using evolutionary algorithms instead of the conventional gradient-based algorithms, it might be perfectly possible for an AI to obtain a survival instinct, even at the current level of AI development. As reviewed earlier (page 40), optimization procedures mimicking evolution are sometimes actually used in AI. Large-scale application of evolutionary methods definitely has the potential of creating a survival instinct in AI agents. It is a necessary logical consequence of fundamental evolutionary pressures:

<sup>7</sup>A highly readable account can be found in *Vanity Fair*, "Elon Musk's Billion-Dollar Crusade to Stop the A.I. Apocalypse", April 26, 2017.



To spread its artificial “genes”, an agent has to survive long enough to produce offspring if the evolutionary optimization method is similar enough to biological evolution.

## Self as desires based on internal rewards

Going back to our main topic, suffering, it is clear that both self-preservation and self-evaluation are important sources of suffering.<sup>8</sup> First, it is well-known that depression and low self-esteem create suffering— and they are largely produced by the self-evaluation system. It is, in fact, rather easy to see this as a form of frustration, so it is very much in line with the ideas of the preceding chapters. Self-evaluation is based on a set standard of how good the self should be, in terms of how much reward it should be able to obtain. If such self-evaluation returns a negative result, that can be seen as a form of frustration, similar to reward loss. One could say that the agent had a long-term *desire* or goal to achieve that standard of average rewards, but the agent failed.

Second, self-preservation is obviously behind (physical) pain, which is signalling when damage is happening to the physical organism, according to the IASP definition of pain (page 15). The same idea was extended to suffering by Cassell’s definition (page 17). He emphasizes “loss of the intactness of person” or “threat” thereof, and that this applies not only to physical intactness but to further aspects such as one’s self-image. Replace his term “person” by “self”, and an interpretation related to the discussion in this chapter is clear: self-preservation mechanisms signalling threats to self—even in a very wide sense of the word—directly create suffering.

Thus, in line with the literature review in Chapter 2, we seem to have two different kinds of suffering related to self-needs. One is born from frustration, in this case based on self-evaluation, and easy to understand by the theories of the preceding chapters. The other kind of suffering comes from a threat to the self, and has only been considered in this chapter. From this viewpoint, we would see suffering related to survival as based on a mechanism which is fundamentally different from frustration, and thus rather different from anything we have treated in the preceding chapters. These two mechanisms might only be similar in the sense that both are forms of error signalling.

However, these seemingly different sources of suffering can be brought together by seeing self-preservation as a form of desire as well. In fact, self-preservation can be seen as a long-term goal or desire which can be frustrated: it is a desire to survive. This is in line with van Hooft’s theory of suffering (page 18), where different aspects of one’s being have different needs, ranging from biological survival to meaning of life. This shows how the two different mechanisms of suffering identified in Chapter 2 have a much closer connection than it might first seem.

In more computational terms, a direct way of linking self and desires is based on defining *internal rewards* (or intrinsic motivation). Reward is, by definition, what an AI agent ultimately wants when it is trained in the conventional framework. As we have seen, just wanting immediate reward is quite short-sighted: If the agent is intelligent enough, it will try to compute the state-value function and thus take future rewards into account.

---

<sup>8</sup>It is actually not my goal here to define what self *is*, I am merely considering how phenomena typically associated with “self” are related to suffering, amplifying or even producing it. In fact, there is some ambiguity in this chapter regarding whether self in, say, self-evaluation is the target of evaluation or the system that evaluates; and whether self-evaluation can be seen as a process that somehow leads to the emergence of self or whether self-evaluation takes some “self” that is already defined and then evaluates its performance. Similar ambiguities hold for self-preservation, as well as the further discussions of self-related phenomena in later chapters, in particular, self as control in Chapter 11 and self-awareness in Chapter 12. It is related to the distinction between the “I” and “me” aspects of self, i.e. self as subject or object, described by William James.

But, even the state-value function framework, with discounted future rewards, may not always provide the best practical solution to the problem of maximizing rewards. This is because the value function may be extremely difficult to learn: there may not be enough data to learn it, and even with enough data, it may be incredibly complex to compute.

Therefore, it has been found that it is often useful to program some additional rewards in the agent, in particular rewards that somehow improve its long-term functioning. That is, the system is programmed to receive internally generated reward signals in addition to actual, “external”, rewards. These internally generated reward signals are treated by the learning and planning systems just as if they were real reward signals. Such internal rewards lead to what is called “intrinsic motivation” for behaviour; it could also be called “intrinsic desire”.

### *Curiosity as an example of internal reward*

As a practical example of such intrinsic reward, let us consider *curiosity*, which is widely used in current AI. The starting point here is that when an agent learns in a real environment, the data it receives is strongly influenced by its own actions. If the robot never enters a room, it will not know what is in that room. The action of deciding to enter or not to enter that room will strongly impact the data it gets about that room. This is a problem since usually, the agent does not know what kind of actions create useful data. Therefore, learning to act intelligently necessarily requires a lot of trial and error. That is, the agent just tries out what happens when you do something rather random in each possible situation. Such exploration is actually imposed on almost any agent learning by reinforcement learning. A very simple way of achieving that is to somehow randomize the actions: for example, in 1% or 10% of the time steps, the agent could take a completely random action just to see what happens.<sup>9</sup>

If you want to buy a new electronic gizmo you have never bought before, a basic exploration strategy would mean you just randomly enter different shops, try to buy it, and depending on whether they sold it to you or not and with what price, you slowly update your value function. Most of your time would probably be spent in trying to buy the gizmo in fashion stores that don’t stock any. Because your actions are quite random, you will end up going to the same stores several times, to the great annoyance of the shop assistants. Since you move around randomly, you easily end up going round and round in the same neighbourhood. Gathering data for reinforcement learning is thus particularly difficult because the agent needs to try out different actions, but if it is done completely randomly, much of the time it will take actions that are not very useful for learning, and don’t bring any reward either.

Here we come to the idea of *curiosity*. It means that the agent does not try out completely random actions, which is very inefficient, but there is an internal mechanism that steers the exploration in an intelligent way. What we are talking about here is designing an intrinsic reward system that leads to particularly intelligent exploration.<sup>10</sup> Basically, the agent should try out new actions if they are *informative*. If the agent has never tried a certain action in a certain state, and it has no information that enables it to infer what such an action would do, it would be useful to just try it out. That is, instead of completely randomly trying out new actions, the agent should try out actions whose effects it does not know and cannot predict. This is a more sophisticated

<sup>9</sup>See e.g. (Sutton and Barto, 2018, Ch. 2); for neuroscience results, (Costa and Averbeck, 2020). Similar randomness may even be useful in the motor system, where it is often, perhaps erroneously, considered unwanted noise (Dhawale et al., 2017).

<sup>10</sup>(Schmidhuber, 1991; Mirolli and Baldassarre, 2013; Pathak et al., 2017; Hazan et al., 2019)

form of exploration, and similar to what we would call curiosity in humans: try out things which you never did before—but don't repeat them once you've seen what happens! An intrinsic reward should then be given to the agent every time it successfully engages in such curious exploration and obtains new information.

Curiosity enables the agent to better learn the general structure of the world it is living in, since it will more systematically explore as many possibilities of action as possible. Such exploration can greatly improve future planning, since the agent will learn a better model of the world, and thus it indirectly contributes to future reward.<sup>11</sup> In the gizmo shopping example above, you would not enter the same store twice, since re-entering the same shop gives little new information. You would actually get an internal reward for going to a different street, even a new neighbourhood, which certainly increases your chances of finding the right kind of store. It is likely that such curiosity has been programmed in animals by evolution.<sup>12</sup>

### *Programming self as internal rewards*

Some aspects of the self could clearly be programmed as internal rewards. Self-preservation is obviously one: Most reasonable programmers would probably assign a large negative reward to the destruction of the agent, since losing the agent is certainly expensive in most cases. Then, the planning system will not go to states leading to the agent being destroyed. In fact, you would ideally program the agent so that it keeps quite far away from anything like destruction. This is possible by programming an internal reward which gives a negative reward at any state that is even close to destruction. In other words, any perceived *threat* triggers a negative internal reward signal.<sup>13</sup> Thus, the agent tries to keep far away from threatening situations, as if it had a desire for self-preservation—or more generally for safety, meaning the absence of threats. This is how we can connect rewards to the IASP definition of pain and Cassell's definition of suffering, where not only damage or loss of intactness causes pain and suffering, but a threat as well ("potential damage" in the IASP definition). If an unexpected threat appears, a negative internal reward (or internal "punishment") is triggered, and that causes reward loss and frustration.<sup>14</sup>

---

<sup>11</sup>An abstract way of justifying curiosity is that basic iterative learning mechanisms such as gradient descent often get stuck in what is called "local minima" of an objective (error) function. That means a point in the parameter space that has a better value of the objective function than any other point near-by, but so that there is a point far-away in the parameter space which has an even better value. A special class of optimization methods called "global optimization" tries to improve iterative algorithms so that they might find the global minimum, that is, the very best value for the parameters, or at least something better than simple gradient descent. Bayesian optimization is one class of such methods (Gutmann et al., 2016; Brochu et al., 2010).

<sup>12</sup>(Singh et al., 2010) One might ask whether such curiosity could not be learned by the agents as part of the reinforcement learning process. That might be possible in principle, but it would probably take too long. An animal would learn to be curious when it has reached a certain age, but it is probably more useful for animals to be curious when they are young, as tends to be the case in biology. In AI, researchers also assume that such curiosity must be explicitly programmed.

<sup>13</sup>A threat has been defined as "anticipation of potential harm" (Palmwood and McBride (2019) based on Folkman and Lazarus). In our formalism, this could tentatively be translated as "prediction of likely (large) reduction in state-value" (or, in a short-sighted calculation, prediction of likely negative reward); the computational definition of threat is an important topic for future research. In the narrow context of self-preservation, we could also talk about the prediction of likely destruction (or death). Our definition is different from what is often given in the literature, where a threat is defined in a social, interpersonal, or legal context, and it is an act made by one agent towards another. In this book (as well as in Cassell's definition of suffering, in my interpretation), threat is not interpersonal in most cases. Yet, any definition of an interpersonal threat could be applied here by assuming that "nature" is the other agent making the threat.

<sup>14</sup>To explicitly apply the definition of reward loss, we would need to define the "expected reward". Clearly, here the expectation is something different from Chapter 5, where it was essentially a prediction (but see footnote 14 in that chapter on discussion of different definitions of the expectation). Here, the expectation is something defined by evolution or the programmer. In fact, the situation is very

As already mentioned, the self-evaluation system is clearly nothing else than an internal reward and punishment system, which steers the agent's behaviour in a certain direction. The difference to ordinary rewards is not only that these self-evaluation rewards come from the internal evaluation system: another fundamental difference is that the self-evaluation system is giving internal rewards to the “learning to learn” system, which sets internal parameters of the system and works on a longer time scale. That system does not directly affect the plans made by the agent, but it tries to improve the general functioning of the planning system to improve all future planning. In animals, such internal rewards are presumably programmed by evolution, since their utility can only be seen in a very long time horizon.<sup>15</sup>

### *Self and suffering in Buddhist philosophy*

In line with the concept of internal rewards, the Buddha mentions three different kinds of desires: desire for sense pleasures, desire to be, and desire not to be. While the first one can be interpreted as desire for rewards in the ordinary AI sense, the “desire to be” can be interpreted as desiring the self to simply be in the sense of surviving, and further that the self should be something particular. These two kinds of “desire to be” correspond to the self-needs as defined in this chapter. (The “desire not to be” is, in this interpretation, the desire that the self is not something which is considered bad.) Thus, even in early Buddhist philosophy, suffering related to self has been to some extent reduced to suffering related to desires and frustration.<sup>16</sup>

In later schools of Buddhism, the importance of self was greatly magnified, and some texts even seem to attribute all desires and all suffering to the existence of the “self” (sometimes translated as the “ego”) or attachment to it. This means viewing the connection between desires and self from the opposite angle, considering the self as the *source* of all desires—instead of the self being the target of some very specific desires as in this chapter. For example, self-preservation requires certain actions to be performed, certain goals to be set, and thus self-preservation leads to desire towards those particular goals.<sup>17</sup>

## **Uncertainty, unpredictability, and uncontrollability as internal frustration**

Another important case of internal rewards concerns the properties of uncertainty, unpredictability, and uncontrollability. They are strongly related to frustration, as has been mentioned earlier, and will be considered in detail in later chapters: if the world is, say, uncontrollable, frustration is difficult to avoid. However, there is another remarkable connection: these properties seem to be sources of suffering in themselves as well.

---

simple if the internal reward system gives a negative reward when a threat appears: the expected reward can be defined as always zero, so any negative reward leads to reward loss. The same logic applies to the self-evaluation system if it operates with similar negative rewards, as discussed next in the text.

<sup>15</sup>The distinction between external and internal rewards, or internal and external motivation as they are called in psychology, may not always be very clear. Both come from the programmer or evolution anyway. See footnote 33 in Chapter 5 which discusses the definition of rewards vs. learned state-values; a similar logic has been applied on internal vs. external rewards by Singh et al. (2010); see also Doya and Uchibe (2005).

<sup>16</sup>See *Samyutta Nikaya* 56.11; my interpretation follows Teasdale and Chaskalson (2011a). Different interpretations are possible: One is that “desire to be” means desire that something in the world should be in a certain way. On the other hand, “desire not to be” could possibly express suicidal tendencies—all these desires were condemned by the Buddha.

<sup>17</sup>Another important way in which frustration and self are related is that frustration is particularly strong if the cause of the frustration is attributed to the self (“it was my fault”). However, such attribution of causes is a complicated issue I will not discuss here; see Mancinelli et al. (2021) for a computational treatment.

One of the most robust findings in studies of economic decision-making is that humans are willing to pay money to reduce uncertainty. Such “risk aversion” can be evolutionarily advantageous and is observed even in animals.<sup>18</sup> In addition to such rational economic calculations, uncertainty also *feels* unpleasant in the body, and it is an important factor in stress. Likewise, lack of control is usually considered detrimental to mental well-being. Psychological experiments show that lack of control and uncertainty even make physical pain feel worse.<sup>19</sup>

We can understand this interplay by using, again, the concept of internal rewards. It could very well be that uncertainty, unpredictability, or uncontrollability are suffering in themselves because they lead to frustration of specific internal rewards. If, say, controllability is lower than some expected standard, a frustration signal could be launched. That would be useful for learning because it signals that the agent has failed in learning about the environment; it should not have got itself into a situation where controllability is that low. This is equivalent to a self-evaluation system which considers that the agent should not be in situations that are uncertain, difficult to predict or difficult to control. By such mechanisms, uncontrollability, as well as uncertainty and unpredictability, can directly lead to suffering.<sup>20</sup>

## Fear, threat, and frustration

Fear is a phenomenon that is central to understanding human suffering. It has an obvious connection to self-needs, in particular survival. In fact, it may seem a bit too abstract to talk about suffering as coming from a survival instinct as I did above: such suffering may always be mediated by a feeling of fear. Fear is actually a multifaceted phenomenon and we will consider various aspects of fear in later chapters (especially Chapter 8). For now, let us just look at fear from the viewpoint of self-needs, and investigate the fundamental mechanism for suffering operating in connection with fear.

Suppose you suddenly find yourself in the presence of a tiger in a jungle. You are likely to suffer at that very moment, but why exactly? It is not that you missed something you wanted to have, or some reward you anticipated, so this is not a case of typical frustration. (Nor is it obviously a case of aversion-based frustration, where you didn’t expect something unpleasant to happen but it did, because the tiger hasn’t yet attacked you.) The case is rather that you are *predicting* something terrible to happen with a non-negligible probability. Aristotle proposed that “Fear may be defined as a pain or disturbance due to a mental picture of some destructive or painful evil *in the future*”<sup>21</sup> (my italics). Such a prediction, or a “mental picture”, contains a threat and falls in the scope of Cassell’s definition of suffering due to threat to the intactness of the self or person.

Just as I discussed in connection with self-preservation, it could be argued that we should consider such suffering produced by fear as fundamentally different from frustration. Yet, we can again sketch an interpre-

<sup>18</sup> (Zhang et al., 2014; Platt and Huettel, 2008)

<sup>19</sup> (Koolhaas et al., 2011; De Berker et al., 2016; Peterson, 1999; Hirsh et al., 2012; Yoshida et al., 2013; Seymour, 2019). Lack of control may increase physical pain because the warning signal in pain has to be taken more seriously when the agent cannot do much about the situation and cannot avoid the threat that causes the warning (Wiech et al., 2008).

<sup>20</sup> Reducing uncertainty as measured by entropy can even be seen as a general learning principle for the brain (Friston, 2010), and thus failure to reduce uncertainty should generate an error signal. At the same time, reducing uncertainty, unpredictability, and uncontrollability is very closely related to the goal of curiosity. Uncertainty can be reduced by a curious investigation of new aspects of the environment, and uncontrollability can be reduced by trying out the effects of actions in new circumstances.

<sup>21</sup> *Rhetoric*, II.5, translated by W. Rhys Roberts. This is often abbreviated as “Fear is pain arising from the anticipation of evil.” Cf. the discussion on the definition of threat in footnote 13 in this chapter.

tation which frames fear as a form of frustration, based on internal rewards. First of all, fear is usually, if not always, accompanied by uncontrollability and unpredictability, which were just seen to produce frustration based on internal rewards and related reward loss. Furthermore, suppose there is another internal reward system based on evaluation of the threats related to different states, as proposed above (page 67): any state which is threatening, i.e. predictive of tissue damage or death, is internally programmed to lead to a negative internal reward. Thus, when entering the state where the tiger unexpectedly appears, such an internal reward system issues a negative reward, and this leads to reward loss, because you didn't expect any such negative reward. This is a simple way of interpreting fear as a frustration: fear is due to the appearance of a threat which leads to internal reward loss.<sup>22</sup>

In this example, the planning system can amplify the frustration, because faced with a threat, planning may be launched with the goal state taken as any state where the threat is not present: You are frantically thinking about what to do to be safe. Planning is attempted, but it fails: No plan is found that would get rid of the threat, or if such a plan is found, its execution fails. Thus, any goal state that would be safe is not reached, and there is frustration even in the sense of plans failing.<sup>23</sup>

### *Fear and the level of intelligence*

A simple AI agent might only generate the suffering signal when something bad happens, such as when it fails in its tasks—this is the basic case of frustration. Suppose a thermostat connected to a heating system tries to keep the room at a constant temperature. (This is actually a task that the nervous systems of many animals face as well.) It continually monitors the room temperature and adjusts its actions accordingly. Its function is based on a simple error signal created when the room gets too hot or too cold. When the temperature is suitable, there would be no error signals whatsoever, and certainly no suffering.

Now, if you make the thermostat very intelligent, so that it is able to predict the future and evaluate itself, perhaps even think about its own survival. Then, it might not only suffer when the room temperature is wrong, but also when it anticipates that that might happen. Your hyperintelligent thermostat might be reading the weather forecast on the internet. Suppose the forecast says tomorrow night will be exceptionally cold, beyond the capacities of the heating system; the thermostat anticipates that tomorrow night it will not be able to keep the temperature high enough. Thus, the thermostat suffers due to such a fear—at least in the computational sense.<sup>24</sup>

<sup>22</sup>This interpretation seems to require that the threat is really about self-preservation, which we have narrowly defined from the viewpoint of physical survival, or the very existence of the agent. However, if we extend the concept of self-preservation and threats as in Cassell's definition of suffering to include preservation of one's self-image (including in the social context), this logic applies more generally. For example, it explains the situation where you are about to give a public talk to a big audience and that makes you scared.

<sup>23</sup>Alternatively, the more sophisticated theory of reinforcement learning may provide another explanation even without such internal rewards. As explained in Chapter 5, in particular footnote 20, the RPE theory says there is frustration solely created by predictions in case you move to a state of lower value, and there is no reward. This could be the case when faced with a tiger if we assume that at every time step your probability of being eaten increases. Then, your chances of getting any positive reward in the future are getting smaller and smaller (because you won't get any after being eaten), and thus the expected total reward during the rest of your life (which is the definition of state-value) is decreasing. The central assumption here is that your chances of being eaten are getting higher and higher in the presence of the tiger; this is clearly true at least in the beginning when you first perceive the tiger. However, it could be argued that such decrease in predicted reward does not really correspond to fear but something like disappointment.

<sup>24</sup>Understanding why and how the thermostat suffers creates an interesting illustration of applying the theoretical ideas of this chapter and the preceding chapters. Let us assume that the thermostat is predicting a frustration of its main goal to be likely to happen in the future. The theoretical problem here is that there is no frustration yet, so it is not obvious how this would lead to suffering now, if we

The extraordinary thing here is that the hyperintelligent thermostat suffers even long before anything bad happens, before, say, actual frustration is produced, merely by virtue of the newly appeared anticipation of negative reinforcement. Becoming more intelligent means the agent can predict bad things, suffer based on those predictions, and thus suffer much more than it did earlier. “One who fears suffering is already suffering from what he fears” according to Michel de Montaigne.<sup>25</sup> Humans suffer enormously because they are too intelligent in this sense, and prone to thinking too much about the future—a theme I will return to in Chapter 9 where I talk about simulation of the future.<sup>26</sup>

Yet, if we humans are so incredibly intelligent, why cannot we just decide not to fear anything? Why cannot we take Montaigne’s point seriously: He suggested—actually talking about his chronic pain due to kidney stones—that there is no point in imagining or anticipating future pain since that simply induces more suffering? This is a complex question where part of the answer is the dual-process nature of human cognition, which will be treated in the following chapters.

---

merely rely on the basic frustration and reward loss theories of Chapters 3 and 5; we need to consider something more. First of all, the interpretation of RPE as being able to compute frustration based on comparing two predictions as in footnote 23 above (or footnote 20 in Chapter 5) is one option: when the thermostat reads the negative weather forecast, its prediction of total future reward will change for the worse, and this triggers frustration. However, in this chapter we found several other explanations for that phenomenon. If the thermostat realizes it is unable to properly control the temperature in the future, the uncontrollability may trigger a negative internal reward, and a reward loss. If this happens often, it could also be that the self-evaluation system of the thermostat concludes that it is not performing its central task well enough, thus leading to frustration due to the self-evaluation. While it may be somewhat contrived to apply the theory of survival mechanisms or threats to the self in this simple example, it is possible that if the thermostat fails to keep the temperature constant, it will be thrown into the garbage bin, and a programmer might decide to explicitly program self-preservation mechanisms that produce a negative internal reward when such “death” seems to be approaching—this is the mechanism of suffering due to threat-based internal reward. Actually, if the thermostat is really hyperintelligent, it might figure out, by itself, this possibility of being plugged out, and try to behave accordingly, thus creating a self-preservation mechanism on its own. There is yet another possible mechanism, not treated in this chapter but in Chapter 9: the prediction of future difficulties could take the form of a simulation that triggers the frustration mechanism based on the frustration present in that simulation. Thus, we see many mechanisms that may create suffering in the future-looking thermostat. Whether these frustration-based mechanisms really corresponds to what one would usually call fear is a complex question that I cannot conclusively answer at this stage.

<sup>25</sup> *Essais*, III, 13

<sup>26</sup> As an aside, let me mention a most amazing interplay of fear and frustration seen in the fear of frustration, arising at the moment of making decisions. A person can be afraid of choosing the wrong flavour for his ice-cream and spend an embarrassingly long time in the decision-making process. His brain may correctly predict that a frustration will happen in the future if it turns out that he does not like the flavour that much after all. Such a fear might be present surprisingly often when humans make decisions (Schwartz, 2004).

## Chapter 7

# Fast and slow intelligence and their problems

This chapter concludes Part I by discussing the connections between the different forms of information processing and frustration we have seen so far. To this end, we need to understand better two different modes of processing in the brain, which coincide with those in modern AI. They were already discussed in Chapter 4: neural networks and logic-based Good Old-Fashioned AI. The idea of two complementary systems, or processes, is ubiquitous in modern neuroscience and psychology. It is assumed that the two systems in the brain work relatively independently of each other while complementing each other's computations. This leads to what is called “dual-process” or “dual-systems” theories. The two systems, or modes of operation, roughly correspond to unconscious processing in the brain's neural networks, and conscious language-based thinking. In this chapter, we go a bit deeper into that distinction due to its great importance for understanding suffering.

Each of the two systems has its own advantages and disadvantages, which is the theme of this chapter, and in fact, a theme to which we will return many times in this book. Neural networks are based on learning: they need a lot of data and result in inflexible functioning. On the other hand, GOFAI relies on well-defined categories which may not be found in reality, and the computations needed may be overwhelming as in planning. On the positive side, we show how the two systems can work together in recent AI systems. We conclude by discussing how the different forms of frustration seen in earlier chapters are related to this two-systems distinction, providing a synthesis of such different forms of suffering.

### Fast and automated vs. slow and deliberative

Let us start by the viewpoint on the two systems given by cognitive psychology and neuroscience.<sup>1</sup> According to such “dual-process” (or “dual-systems”) theories, one of the two systems in the brain is similar to the neural networks in AI: It performs its computation very fast, and in an automated manner. It is fast thanks to its computation being massively parallel, i.e., happening in many tiny “processors” at the same time. It is automated in the sense that the computations are performed without any conscious decision to do so, and without any feeling of effort. If visual input comes to your eyes, it will be processed without your deciding to do so, and usually you recognize a cat or a dog in your visual field right away, that is, in something like one-tenth of a second.<sup>2</sup> Most of the processing in this system is also unconscious. You don't even understand how the computations

---

<sup>1</sup>(Evans, 2008; Kahneman, 2011; Sloman, 1996)

<sup>2</sup>(Kirchner and Thorpe, 2006)



are made; the result of, say, visual recognition just somehow appears in your mind, which is why this system is also called “implicit”.

The processing in the conscious, GOFAI-like system is very different. To begin with, it is much slower. Consider planning how to get home from a restaurant where you are the first time: you can easily spent several seconds, even minutes, solving this planning task. The main reason is that the computations are not parallelized: They work in a serial way, one command by another, so the speed is limited by the speed of a single processing unit. In humans, another reason why symbolic processing is slow is, presumably, that it is evolutionarily a very new system, and thus not very well optimized. Other typical features of such processing are that you need to concentrate on solving the problem, the processing takes some mental effort, and it can make you tired. Such processing is also usually conscious, which means that you can explain how you arrived at your conclusion; hence the system is also called “explicit”.<sup>3</sup>

Note that in an ordinary computer, the situation above is in some ways reversed, as already explained on page 43 in Chapter 4. A computer can do logical operations much faster than neural network computations, since logical operations are in line with its internal architecture. In fact, a computer can only do neural network computations based on a rather cumbersome conversion of such analog operations into logical ones. Presumably, in fact, the brain is only able to do logical operations after converting them into neural network computations, and that is equally cumbersome.

To see the division into two systems particularly clearly, we can consider situations where the two systems try to accomplish the *same* task, say, classification of visual input. We can have a neural network that proposes a solution, as well as a logic-based system that proposes its own. Sometimes, the systems may agree; at other times, they disagree.

Suppose a cat enters your visual field. When the conditions for object recognition are good, your visual neural network would recognize it as a cat. In other words, the network would output the classification “cat” with high certainty. However, when it is dark, and you only get a faint glimpse of the cat that runs behind some bushes, your neural network might not be able to resolve the categorization. It might say it is probably either a cat or a dog, but it cannot say which. At this point, the more conscious, logic-based system might take over. You recall that your neighbour has a cat; you don’t know anybody who owns a dog near-by; you think this is just the right moment in the evening for a cat to hunt for mice. Thus, you logically conclude it was probably a cat. In this case, the task of recognizing an object used the two different systems, working together. The logic-based one took quite some time and effort to use, while the neural network gave its output immediately and without any effort. Here, the systems were not completely independent, since the logic-based system did need input from the neural network to have some options to work on.

The two systems can also disagree, as often happens in the case of fear. Talking about fear and related emotional reactions, people often call them “hard-wired”. This expression is not too far from reality. What happens is that the brain uses special shortcut connections to relay information from the eye to a region called the amygdala, an emotional center in the brain. This shortcut by-passes those areas where visual information is usually processed.<sup>4</sup> If such a connection learns to elicit fear (due to a previous unpleasant encounter with

<sup>3</sup>My exposition is a kind of synthesis of different theories, and not all the mentioned properties are always associated with the two systems. Further, I should mention the proposals that the second, explicit system may be specialized in simulating hypothetical events that have not happened (Stanovich, 2004) for example for the purposes of planning, which will be considered in Chapter 9; or it could be mainly about working memory (Evans, 2008). An interesting related division between feedforward and feedback processing in neural networks is discussed by Lamme and Roelfsema (2000).

<sup>4</sup>(LeDoux and Pine, 2016). More precisely, a pathway goes directly from the thalamus to the amygdala, without reaching the visual

some animals, for example), it will be very difficult to get rid of it. Any amount of reasoning is futile, presumably since the visual signal triggering fear is processed by completely different brain areas than logical, conceptual reasoning. Often, the logic-based system loses here, and the neural-network-based fear prevails. This division into two processes also explains why it is difficult for us to change unconscious associations, such as fear: the conscious, symbolic processing has limited power over the neural networks.

Interestingly, people tend to think that the main information processing in our brain happens by the conscious, symbolic system, including our internal speech and conceptual thinking. But what if that is simply the tip of the iceberg, as early psychoanalysts<sup>5</sup> claimed more than a hundred years ago? The idea that most information processing is conscious and conceptual may very well be an illusion. We may have such an impression because conceptual processing requires more effort, or because it is more accessible to us by virtue of being conscious. However, if you quantify the amount of computational resources which are used for conceptual, logical thinking, and compare them with those used for, say, vision, it is surely vision that will be the winner.<sup>6</sup>

Similar to the dual-process theories in cognitive psychology and neuroscience just described, the division between GOFAI and neural networks has been prominent in the history of AI research, which has largely oscillated between the two paradigms. Currently, neural networks are very popular, while GOFAI is not used very widely. However, this may very well change, and perhaps in the future, AI will combine logic-based and neural models in a balanced way. Since GOFAI is used by humans, it is very likely to have some distinct advantage over neural networks, at least for some tasks.<sup>7</sup>

Note that in AI we find another important distinction which is not very prominent in the neuroscientific literature: learning vs. no learning. Neural networks in AI are fundamentally based on learning, and using them without learning is not feasible. In contrast, in its original form, Good Old-Fashioned AI promises to deliver intelligence without any learning. That comes at the cost of much more computation, and more efforts spent on programming. This distinction seems to be relevant to the brain, even if not as strict as in AI, as we will see next.<sup>8</sup>

## Neural network learning is slow, data-hungry, and inflexible

To understand the relative advantages of the two systems, let us first consider the limitations in neural networks, and especially the learning that they depend on. First of all, neural network learning is data-hungry: it needs large amounts of data. This is because the learning is by its very nature *statistical*; that is, it learns based on statistical regularities, such as correlations. Computing any statistical regularities necessarily needs a lot of data; you cannot compute statistics by just observing, say, two or three numbers.

---

cortex. Another, slower pathway does go back from the visual cortex to the amygdala. See also Hofmann et al. (2009) on conflicts between the two systems from the viewpoint of self-control.

<sup>5</sup>I'm here obviously referring to Freud and his followers, but the importance of unconscious processing was emphasized around the same time frame by Janet (1889), and even earlier by philosophers such as Arthur Schopenhauer and Eduard von Hartmann.

<sup>6</sup>(Nakayama, 1999). While this comparison in terms of brain resources seems compelling in terms of comparing vision vs. conceptual thinking, it is more difficult to compare the conscious and unconscious aspects since we don't really know how consciousness is related to the brain; see Chapter 12.

<sup>7</sup>While the main text discusses later some such combinations of the two systems, I should also mention attempts made under the titles of "hybrid AI" or "neural-symbolic processing" (d'Avila Garcez et al., 2012; Goertzel, 2012; Graves et al., 2016; Yi et al., 2018).

<sup>8</sup>When we talk about learning in the brain in the context of neural network models, that is to be understood on an abstract level, where learning includes both evolutionary and developmental processes; this will be discussed in more detail in Chapter 10, page 115.

Second, neural network learning is slow. Often, it is based on gradient optimization, which is *iterative*, and needs a lot of such iterations. The same applies to Hebbian learning, where changing neural connections takes many repetitions of the input-output pairs—this is natural since Hebbian learning can be seen as a special case of stochastic gradient descent. In fact, to input a really large number of data points into a learning system almost necessarily requires a lot of computation, since each data point takes some small amount of time to process.

This statistical and iterative nature of neural network learning has wide-ranging implications for AI. To begin with, these properties help us to further understand why it is so difficult, in us humans, to change any kind of deeply ingrained associations. Mental associations are presumably in a rather tight correspondence with neural connections: If you associate X with Y, it is because there are physical neural connections between the neurons representing X and Y. Now, even if any statistical connection ceases to exist in the real world, perhaps because you move to live in a new environment, it will take a long time before the Hebbian mechanisms learn to remove the association between X and Y, or to associate X with something else.<sup>9</sup>

In fact, these learning rules, whether basic Hebbian learning or some other stochastic gradient methods, may seem rather inadequate as an explanation for human learning: We humans can learn from single examples and do not always need a lot of data. You only need to hear somebody say once “Helsinki is the capital of Finland”, and you have learned it, at least for a while. Surely, you don’t need to hear it one thousand times, although that may help. This does not invalidate the neural network models, however, since the brain has multiple memory systems, and Hebbian learning is only one way we learn things and remember them—we will get back to this point in Chapter 9.<sup>10</sup>

The iterative nature of neural learning, together with the two-process theory, also helps to explain in more detail why it is so difficult to *deliberately* change unconscious associations. Suppose you consciously decide to learn an unconscious association between X and Y (where X might be “exercise” and Y might be “good”). How can you transfer such information from the conscious, explicit system to the neural networks? Perhaps the best you can do is to recall X and Y simultaneously to your mind—but that has to be done many times! In fact, you are kind of creating a kind of new data and feeding it into the unconscious association learning in your brain. You are almost cheating your brain by pretending that you perceive the association “X and Y” many times. We will see many variations on this technique when we consider methods for reducing suffering in Chapter 15.

Another limitation is that when a neural network learns something, it is strictly based on the specific input and output it has been trained on. While this may seem like an obvious and innocuous property, it is actually another major limitation of modern AI. Suppose that a neural network in a robot is trained to recognize animals of different species: It can tell if a picture depicts a cat or a dog, or any other species in the training set. Next, suppose somebody just replaces the camera in the robot with a new one, with higher resolution. What happens is that the neural network the robot previously trained does not work anymore. It will have no idea how to interpret the high-resolution images since they do not match the templates it learned for the original data. A

---

<sup>9</sup>In some cases, an association may not actually be removed but overridden by an inhibitory connection, a bit like creating a new “negative” connection to cancel the functioning of a positive connection (Westbrook et al., 2002). This also means the old association can be reactivated quite easily.

<sup>10</sup>Chapter 9 will explain the idea of *replay* whose application to this case would be as follows. Maybe your brain does actually hear the sentence “Helsinki is the capital of Finland” many times. One of the learning systems in the brain is based on storing events, or short episodes, in an area called the hippocampus. It uses special mechanisms, presumably quite different from stochastic gradient methods, to store the sentence after hearing it just once. Then, the hippocampus feeds the sentence to the other parts of the brain many times, and that allows Hebbian learning and something similar to stochastic gradient learning to happen.

similar problem is that the learning is dependent on the context: An AI might be trained by images where cats tend to be indoors and dogs outdoors, and it will then erroneously classify any animal pictured indoors as a cat. The AI sees a strong correlation between the surroundings and the animal species, and it will not understand that the actual task is about recognizing the animals and not recognizing the surroundings. That is why a neural network will typically only work in the environment or context it is trained in.<sup>11</sup>

In light of these limitations, AI based on neural networks is thus rather different from what intelligence usually is supposed to be like in humans. In a celebrated experiment, human participants started wearing goggles containing a prism which made their world look upside down. Surprisingly soon, the participants were able to function normally; somehow, their visual systems were able to process the input correctly in spite of the inverted visual input.<sup>12</sup> In general, when humans learn to perform a task, they are often somehow able to abstract general knowledge out of the learning material, and they are able to transfer such knowledge from one task to another. It has even been argued that the hallmark of real intelligence is that it is able to function in many different kinds of environments and accomplish a variety of tasks without having to learn everything from scratch. If all a robot can do is to mow the lawn, we would think it is just accomplishing a mechanical task and is not “really” intelligent.<sup>13</sup>

## Using planning and habits together

Let us next look at how the two systems might interact in AI. Regarding action selection, we have actually seen how two different approaches can solve the same problem in AI: reinforcement learning and planning. Planning is in fact one of the core ideas of the GOF AI theory. Planning is undeniably a highly sophisticated and demanding computational activity, and probably impossible for simple animals—some would even claim it is only present in humans, although that is a hotly debated question.<sup>14</sup> In any case, it seems to correspond closely to the view humans have about their own intelligence, and therefore was the target of early AI research. However, in the 1980s, there was growing recognition that building agents, perhaps robots, whose actions show human-level intelligence is extremely difficult, and it may be better to set the ambitions lower. Perhaps building a robot which has the level of intelligence of some simple animal would be a more realistic goal. Moreover, like in other fields of AI, learning gained prominence. That is why habit-like reinforcement learning started to be seen as an interesting alternative to planning.<sup>15</sup>

### *Habits die hard—and are hard to learn*

However, habit-based behaviour has its problems, partly similar to those considered above for neural network learning. Learning the value function, that is, learning habits, obeys the same laws as other kinds of machine learning. It needs a lot of data: the agent needs to go and act in the world many, many times. This is a major bottleneck in teaching AI and robots to behave intelligently, since it may take a lot of time and energy to make,

---

<sup>11</sup> (Arjovsky et al., 2019)

<sup>12</sup> For a recent review, see (Pisella et al., 2006).

<sup>13</sup> (Legg and Hutter, 2007). Functioning in many environments thus requires an advanced capacity to what is called *transfer learning*, which is currently a focus of very active research in AI (Pan and Yang, 2009; Weiss et al., 2016).

<sup>14</sup> (Redshaw and Bulley, 2018; Corballis, 2019)

<sup>15</sup> A related school of research emphasized how intelligence might emerge from simple reactive behaviours, even without any learning (Brooks, 1991, 1999).

say, a cleaning robot try to clean the room thousands of times. Basic reinforcement algorithms are also similar to neural network algorithms in that they work by adjusting parameters in the system little by little, based on something like the stochastic gradient methods.

Another limitation which is crucial here is that the result of the learning, the state- or action-value function, is very context-specific. If the robot has learned the value function for cleaning a room, it may not work when it has to clean a garden. Even different rooms to clean may require slightly different value functions! The world could also change. Suppose the fridge from which the robot fetches the orange juice for its master is next to a red closet. Then, the robot will associate the red closet with high value since seeing it, the robot knows it is close to being able to get the juice. However, if somebody moves the closet to a different room, the robot will start acting in a seemingly very stupid way: It will go to the room which now has the red closet when it is supposed to get the orange juice—in fact, it might simply approach any new red object introduced to its environment in the hope that this is how it finds the fridge. It will need to re-learn its action-values all over again.

Here we come to the other side of the slowness of learning habits: Once a habit is learned, it is difficult to get rid of it. In humans, the system learning and computing the reinforcement value function is outside of any conscious control: We cannot tell it to associate a smaller or larger value to some event. This is why we often do things we would prefer not to do, out of habit. In order to learn that a habit is pointless in the sense that it does not give any reward anymore (as happened with the robot above), a new learning process has to happen, and this is just as slow as the initial learning of the habit. That is why habits die hard.<sup>16</sup>

### *Combining habits and planning*

These problems motivate a recent trend in AI: combining planning and habit-like behaviour. The habit-based framework using reinforcement learning will lead to fast but inflexible action selection, and is ideally complemented by a planning mechanism which searches an action tree a few steps ahead—as many as computationally possible. Depending on the circumstances, the action recommended by either of the two systems can then be implemented.<sup>17</sup>

Let us go back to the robot which is trying to get the orange juice from the fridge. One possible way of implementing a combination of planning and habit-like behaviour is to have a habit-based system *help* the planning system in the *tree search*. Using reinforcement learning, you could train a habit-based system so that when the robot is in front of the fridge whose door is closed, the system suggests the action “open the door”. When the door of the fridge is open with orange juice inside, the habit-based system suggests “grab the orange juice”. While these outputs could be directly used for selecting actions, the point here is that we can use them as mere suggestions to a planning system. Such suggestions would greatly facilitate planning: The search can concentrate on those paths which start with the action suggested by the habit-based system, focusing the search and reducing its complexity. However, the planning system would still be able to correct any errors in the habit-like system, and could override it if the habit turns out to be completely inadequate.

One very successful real-world application using such a dual-process approach is AlphaGo, a system playing the board game of Go better than any human player.<sup>18</sup> The tree to be searched in planning consists of moves by the AI and its opponent. This is a classical planning problem in a GOF AI sense. The world has a finite

<sup>16</sup>However, some hope will be offered in Chapter 9 where we consider ways of speeding up learning by replaying existing data, and that theme is continued in Chapter 15.

<sup>17</sup>(Daw et al., 2005)

<sup>18</sup>(Silver et al., 2016)

number of well-defined states, and also, the actions and their effects on the world are clearly defined, based on the rules of the game. What is a bit different is that there is an opponent whose actions are unpredictable; however, that is not a big problem because the agent can assume that the opponent chooses its actions using the same planning engine the agent uses itself.

The search tree in Go is huge since the number of possible moves at any given point of the game is quite large, even larger than in chess. In fact, the number of possible board positions (positions of all the stones on the board) is larger than the number of atoms in the universe—highlighting the fundamental problem in GOFAI-style planning. Since it is computationally impossible to exhaustively search the whole tree, AlphaGo randomly tries out as many paths as it has time for. This leads to a “randomized” tree search method called Monte Carlo Tree Search. Algorithms having some randomness deliberately programmed in them are often called Monte Carlo methods after the name of a famous casino. However, even such a random search is obviously quite slow and unreliable.<sup>19</sup>

The crucial ingredient in AlphaGo is another system which learns habit-like behaviours. This system is used inside the planning system, a bit like in the juice robot just described. While the system is rather complex, let's just consider the fact that in the initial stage of the learning, AlphaGo looks at a large database of games played by human experts. Using that data, it trains a neural network to predict what human experts would do in a given board position—the board positions correspond to the states here. The neural network is very similar to those used in computer vision, and gets as input a visual view of the Go board. This part of the action selection system could be interpreted as learning a “habit”, i.e., an instinctual way of playing the game without any planning.<sup>20</sup> The action proposed by the habit system can be used as such, but even more intelligent performance is obtained by using it as a heuristic for the tree search: the tree search is focused on paths related to that proposed action. This heuristic is further refined by further learning stages. In particular, the system also learns to approximate the state-values by another neural network.<sup>21</sup>

Such suggestions based on neural networks are fast, and intuitively similar to what humans would do. Often, a single glimpse at the scene in front of your eyes will tell a lot about where reward can be obtained, and suggests what you should do. Even when humans are engaged in planning, such input coming from neural networks often guides the planning. If you go to get something from the fridge, don't you have almost automated reactions to seeing the fridge door closed, and seeing your favourite food or drink inside the fridge? These are presumably given by a simple neural network. Yet, there is a deliberative, thinking aspect in your behaviour, and you can change it if you realize, for example, that the juice has gone bad—which the simple neural network did not know.

What is typical in humans is that action selection can also switch from one system to another as a function of practice. Learning a new skill, such as driving a car, is a good example—skills are similar to habits from the

<sup>19</sup>(Browne et al., 2012; Chaslot et al., 2008). Monte Carlo Tree Search does include clever tricks which make the search a bit more intelligent. It does not try out actions (or moves in a game) completely randomly, but gathers data on which actions look more promising. In particular, there is quite a lot of data regarding actions taken in the first steps of the search path, since any search has to always try out one of those, and their number is limited because there has not yet been a combinatorial explosion as in the number of long paths. Monte Carlo Tree Search uses such data to bias the search towards paths whose initial parts have been found the most promising.

<sup>20</sup>Interestingly, the “habits” are here learned based on imitation since they are simply trying to replicate what the human players did earlier. Imitation learning is another principle for machine learning, especially important for robots (Schaal, 1999).

<sup>21</sup>For the general theory on approximating values by neural networks or simpler methods, see Sutton and Barto (2018, Chapter 9). In Chapter 9 we will also see how the system can improve by playing against itself. A completely different purpose for combining learning and planning is to learn to plan better in a given environment where rewards are changing (Tamar et al., 2016; Pascanu et al., 2017).

computational viewpoint. First, you really have to concentrate and consciously think about different action possibilities. With increasing practice and learning, you need to think less and less, since something like a value function is being formed in the brain. In the end, your actions become highly *automated*, and you don't really need to think about what you are doing anymore. The habit-based system takes over and drives the car effortlessly.<sup>22</sup>

## Advantages of categories and symbols

While in this example of Go playing, neural networks and GOFAI work nicely together, it is actually not easy in most other tasks to show any clear utility of symbolic AI approaches. This may of course change any time, since AI is a field of rapid development. It is quite likely that GOFAI is necessary for particularly advanced intelligence—something much more advanced than what we have at this moment. Yet, the tendency has recently been almost the opposite: Tasks which were previously thought to be particularly suitable for symbolic AI have been more successfully solved by neural approaches.<sup>23</sup> Perhaps symbolic AI works with board games only because such games are in a sense discrete-valued: the stones on the Go board can only be in a limited number of positions, so the game is inherently suitable for GOFAI. So, we have to think hard about what might be the general advantages of logic-based intelligence compared to neural networks. In the following, I explore some possibilities.

### *GOFAI is more flexible and facilitates generalization*

Suppose that there is a neural network that recognizes objects in the world and outputs the category of each object. Then, what would be the utility of operating on those categories as discrete entities, using symbolic-logical processing, instead of having just a huge neural network that does all the processing needed?

We have already seen, more than once, one great promise of GOFAI in the case of planning: flexibility. Given any current state and any goal state, a planning system can, if the computational resources are sufficient, find a plan to get there. If anything changes in the environment—say, it is no longer possible to transition between two states due to some kind of blockage—the planning system takes that into account without any problems. This is in contrast to reinforcement learning which will not know what to do if the environment changes; it may have to spend a lot of time re-learning its value functions.

Furthermore, GOFAI is easily capable of representing various kinds of *data structures* and relationships in the same way as a computer database. For example, it can easily represent the fact that both cats and dogs are animals, i.e. the hierarchical structure of the categories. It can also represent the relationship that the character string “Scooby” is the name of a particular dog. This adds to the flexibility of GOFAI by allowing more abstract kinds of processing, which are easily performed by humans.<sup>24</sup>

Even without going into logical operations, we can consider the advantages of using discrete categories (cats vs. dogs) instead of a high-dimensional feature space with continuous values. A wide-spread idea is that categories are useful for generalizing knowledge over categories, which in its turn underlies various forms of

<sup>22</sup>I may be oversimplifying things here, since in the brain, learning motor skills such as driving is not quite the same as forming habits, and they may be based on different brain systems. However, on a more abstract level where we only consider the computational principles, they can be very similar (Doyon et al., 2003; Peters et al., 2011; Sun et al., 2005).

<sup>23</sup>Natural language processing is a good example (Mikolov et al., 2010; Bengio et al., 2003).

<sup>24</sup>For an attempt to do these in a neural system, see Frady et al. (2020).

abstract thinking. Even though cats are not all the same, it is useful to learn some of their general properties. They like milk, they purr; they don't like to chew bones like dogs do, and they are not dangerous like bears. Having categories enables the system to learn to associate various properties to the whole category: Observing a few cats drink milk, the system learns to associate milk-drinking to the whole category of cats, instead of just some individual cats. Importantly, associating properties to categories means the system was able to *generalize*: after seeing *some* of the cats drink milk, it inferred that *all* cats drink milk. Such generalization is clearly an important part of intelligence. If the system needed to learn such a property separately for each cat, it would be in great trouble when it sees a new cat and needs to feed it — it would have no idea what to do. But, learning that the whole category of cats is associated with milk-drinking, it knows, immediately and without any further data, what to give to this new cat.

### *Categories enable communication*

Nevertheless, I think the feature which makes GOFAI fundamentally different from neural networks is that the use of symbols is similar to using some kind of a primitive *language*. In fact, you can hardly have GOFAI without some kind of a language—perhaps akin to a programming language—in which the symbols and logical rules are expressed. It is equally clear that with humans, language is primarily used for communication between individuals. As each category typically corresponds to a word, humans can communicate associations, or properties of categories, to each other. I can tell my friend that cats drink milk, so she does not need to learn what to feed to cats by trial and error. I have condensed my extensive data on cats' eating habits into a short verbal message that I transmit to her.

So, it is plausible that the main reason humans are capable of symbolic thinking is that it enables them to communicate with each other. After such a communication system was developed during evolution, humans then started using the same system for various kinds of intelligent processing even when alone. Perhaps we started by telling others, for example, where to find prey.<sup>25</sup> This led to the development of symbols and logical operations, which were found useful for abstract thinking: Perhaps you could try to figure out yourself where you should hunt tomorrow. Eventually, such capabilities ended up producing things such as in quantum physics—and the very theory of GOFAI.<sup>26</sup>

A reflection of the utility of categories in communication may be seen in a recent research line in AI which tries to develop systems whose function is easy to interpret by humans.<sup>27</sup> If you use a neural network to recognize a pattern, the output may be clear and comprehensible, but the computations—why did the network give that particular output—are extremely difficult to understand for humans. This is fine in many cases, but sometimes it is necessary to explain the decision to humans. For example, if an AI rejects your loan application, the bank using the AI may be legally obliged to explain the grounds for that decision.<sup>28</sup> Researchers developing

---

<sup>25</sup> (Wagner et al., 2003)

<sup>26</sup> This is what is called “exaptation” in evolution. It means that a trait was first produced to adapt for one phenomenon, but then it turned out to be useful for something else. A typical example is bird's feathers, which probably first evolved to keep the birds warm, and only later turned out to be useful for flying.

<sup>27</sup> (Zeng et al., 2017; Su et al., 2015; Tan et al., 2017; Guidotti et al., 2018)

<sup>28</sup> For example, the General Data Protection Regulation (GDPR) of the European Union imposes a general “right for explanation” for almost any decision made by an algorithm on an individual (Goodman and Flaxman, 2017). One reason for such a requirement is to make sure that the AI did not discriminate applicants based on gender, race, or similar characteristics—an objective called “fair AI”. Another reason is that the AI might not make the final decision, but could be used as a support for a human decision-maker, such as a medical doctor: The human decision-maker would greatly benefit from understanding why the AI came to its conclusion. One more



such interpretable AI usually end up doing something similar to GOFAI boosted by learning, since it gives rules which can be expressed in more or less ordinary language, and thus they can be explained. In fact, in Chapter 4 we saw examples of GOFAI systems whose functioning is easy to understand and to explain.<sup>29</sup>

## Categorization is fuzzy, uncertain, and arbitrary

Now, let us consider the flipside: problems that arise when using categories. We have already seen some problems in the logical-symbolic processing, the most typical being the exponential explosion of computation in planning. Here, we focus on the consequences of using categories, and look at the question from a more philosophical angle. Indeed, it has been widely recognized by philosophers over the centuries that dividing the world into “crisp” categories can only be an approximation of the overwhelming complexity of the world. I focus on some issues which will in later chapters be seen to be relevant for suffering.<sup>30</sup>

### *Categories are fuzzy*

Philosophers have long pointed out that there may not be any clearly defined categories in the world. Granted, the difference between cats and dogs may be rather clear, but what about the category of, say, a “game”? Wittgenstein gave this as an example of a category which has no clear boundaries. Different games have just some vague similarity, which he called “family resemblance”.

This idea has been very influential in AI under the heading of *fuzziness*. A category is called fuzzy if its boundaries are not clear or well-defined. Consider for example the word “big”. How does one define the category of big things? For simplicity, let us just consider the context of cities. If we say “London is big”, that is clearly true: London definitely belongs to the category of big things, in particular big cities. But if we say “Brussels is big”, is that true or false? How does one define what is big and what is not? In the case of cities, we could define a threshold for the population, but how would we decide what it should be? An AI might learn to categorize cities into big and small ones based on some classification task—in Chapter 4, we discussed how this

---

reason for making AI easy to interpret is that understanding how an AI works makes it easier to evaluate its potential safety hazards, and develop AI that is safe.

<sup>29</sup>I am actually tempted to think that the *only* specific utility of categories (which cannot be obtained without them) is communication, including being interpretable and comprehensible by humans in the case of AI. In particular, it is not clear to me if explicit categories are needed for generalization. Without going into details, let me just mention that similar operations could be performed directly in a representational space by simply propagating any associations to near-by points in that space without any strict division into categories. For opposite viewpoints putting concepts at the heart of (human) cognition, see Rosch (1999); Harnad (2017). Obviously, it is important here to compare the different definitions of categorization used: Harnad (2017) uses a definition which is very general.

<sup>30</sup>I don't go into any details on how that “division” of the world into categories happens, but for the interested reader, I give some pointers here. Earlier, we considered the case where the neural network recognizes an object and outputs its category. This is a simple starting point; while it can be easily done by supervised learning, it can also be implemented by unsupervised learning methods, in particular methods such as clustering and (Gaussian) mixture modelling. In the case of humans, the connection between neural networks and logic-symbolic processing is related to what is called the symbol grounding problem (Harnad, 1990). It is a topic subject to a lot of debate: some argue no proposed solution is sufficient (Taddeo and Floridi, 2005), while others argue it is essential to consider robots which communicate with each other (Steels, 2008). The operation of neural networks is closely related to one well-known proposal called the prototype theory. It means we define each category by a single point in the space the activities of units in a neural network (preferably in layers close to output); this point is the prototype (Rosch, 1978). Basically, you would find a “prototypical” cat as a point in the very center of all those points that represent cats. A generalization of this idea can be found in Gärdenfors (2004). However, things get much more complicated in the case of abstract categories such as “good” or “beautiful”.

might happen in categorizing body temperature into “high fever” or not. However, that categorization would depend on the task, and there would always be a grey zone where the division is rather arbitrary.

The consensus in AI research is that many categories are quite fuzzy and have no clear boundaries; there are only different degrees of membership to a category. There is no way of defining a word like “big” (or, say, “nice”, “tall”, “funny”) in a purely binary fashion. There will always be objects that quite clearly belong to the category and objects which clearly do not belong to the category, but for a lot of objects the situation is not clear. In the theory of fuzzy logic, such fuzziness is modelled by giving each object a number between 0 and 1 to express the degrees of membership to each category.<sup>31</sup>

### *Categorization is uncertain*

In addition, categorization is always more or less uncertain. Any information gleaned from incoming sensory input is uncertain, for reasons we will consider in more detail in Chapter 10. Partly, it is a question of the neural network getting limited information, and partly because of its limited information-processing capabilities. If you have a photograph of a cat taken in the dark and from a bad angle, the neural network or indeed any human observer may not be sure about what it is. They might say it is a cat with 60% probability, but it could be something else as well. In other words, any categorization by an AI is very often a matter of probabilities.

It is important to understand that fuzziness and uncertainty are two very different things. Uncertainty is a question of probabilities, and probabilities are about lack of information. If I say that a coin flip is heads with 50% probability and tails with 50% probability, there is no fuzziness about which one it is. After flipping the coin I can say if it is heads or tails, and no reasonable observer would disagree with me (except in some very, very rare cases). In other words, uncertainty is a question of not knowing what will happen or has happened, i.e., a lack of information about the world. In contrast, fuzziness has nothing to do with lack of information; it is about the lack of clear definition. We cannot say if the statement “Brussels is big” is true even if we have every possible piece of information about Brussels, including its exact population count. According to the information I find on Wikipedia, its population is 1,191,604, but knowing that will not help me with the problem if I don’t know how many inhabitants are required for a city to be in the “big” category.

Humans are not good at processing uncertainty. Various experiments show that humans tend to use excessively *categorical thinking*, where the uncertainty about the category membership is neglected. That is, when you see something which looks to you most probably like a cat, your cognitive system tends to ignore any other possibilities, and think it is a cat for sure.<sup>32</sup>

An old Buddhist parable about these dangers in categorization is seeing a rope in the dark and thinking it is a snake. You miscategorize the rope, and your brain activates not only the category of a snake, but all the associations related to that category (“animal”, “dangerous”). You get scared, with all the included physiological changes, such as an increased heart rate. If you had properly taken uncertainty of such categorization into

<sup>31</sup> (Mendel, 1995)

<sup>32</sup> An example was seen in a study where the subjects were told a story which suggested that an imaginary person entering a house would be either a burglar, or a real estate agent. When the imaginary person was more likely to be a real estate agent than a burglar—based on various cues such as what other characters in the story were thinking—they tended to ignore the possibility that the person is a burglar altogether, as seen in the predictions that they made about the behaviour of the imaginary person (Malt et al., 1995; Murphy and Ross, 2010). The authors also found a way of remedying the situation: if the subjects are asked what the probabilities of the two categories are, and their estimates are shown on the computer screen (say, “65% vs. 35%”), the subjects are able to take the uncertainty of the categorization into account.

account, your reaction might have been more moderate.

### *Categorization is arbitrary*

In some cases, the categories are not just fuzzy or uncertain: Their very existence can be questionable. Consider concepts such as “freedom” or “good”. Even forgetting about any difficulties in programming an AI to understand them, is it even clear what these words mean? Certainly, they mean different things to different people: people from different cultural backgrounds may easily misunderstand each other simply because they use such concepts with slightly different meanings. A great amount of time can be spent in attempting to just describe the meanings of certain words and categories. In fact, we spend more than one chapter on analysing the category called “self” in this book.

Even in rather straightforward biomedical applications of machine learning, we often use categories that are not well-defined. For example, in a medical diagnosis context, it is not clear if what we usually call schizophrenia is a single disease. Perhaps there are a number of different diseases which all lead to the single diagnosis of schizophrenia.<sup>33</sup> Developing effective medications may only be possible once we understand all the subtypes, while thinking of all the subtypes as a single disease (a single category) may mislead any treatment attempts.

Moreover, a categorization that works for one purpose might not be suitable for another. We might divide people into different nationalities, which is very useful from the viewpoint of knowing what languages they are likely to understand. However, we can too easily use the same categories to predict all kinds of personality traits of those individuals, and that prediction may go quite wrong. In more general terms, categories and their utility depend on the context; different people use categories in different ways, thus they are subjective.

Such arbitrariness of categories has been well appreciated in some philosophical schools. In the Yogācāra school of Buddhism, it is claimed that “while such objects [as chairs and trees] are admissible as conventions, in more precise terms there are no chairs, or trees. These are merely words and concepts by which we gather and interpret discrete sensations that arise moment by moment in a causal flux.”<sup>34</sup> What arises in such a moment-by-moment flux is, in our terminology, activities in neural networks. Categories are created afterwards, by further information-processing.

### *Overgeneralization*

It may be easy to understand that miscategorization leads to problems, as in mistaking a rope for a snake. However, the biggest computational problem caused by all properties just discussed—fuzziness, uncertainty, arbitrariness—may be overgeneralization. Overgeneralization can be difficult to spot, even after the fact, which makes it particularly treacherous.

Overgeneralization means that you consider all instances of a category to have certain properties, even if those properties hold only for some of them. Since categories are fuzzy, anything which is not really firmly inside the category may actually be quite different from its prototype. Related to this, you may not acknowledge

<sup>33</sup>(Peralta and Cuesta, 2001; Brodersen et al., 2014). The same could be said of depression (Drysdale et al., 2017). For a general overview on such “precision medicine”, see Insel and Cuthbert (2015).

<sup>34</sup>Quote from (Lusthaus, 2013), see also (Lusthaus, 1998; Tagawa, 2009; Williams, 2008b). In ancient Greece, Pyrrhonian Skeptics had ideas similar to such Buddhist schools, and indeed Pyrrhonians are likely to have been directly influenced by Buddhist thinkers due to Alexander the Great’s campaign into India (Bowie, 2016; McEvilley, 1982; Garfield, 1990).

the uncertainty of categorization and the ensuing generalization. Even more rarely do people acknowledge that the very categories are arbitrary.

Overgeneralization effects are well documented, for example, in perception of human faces, where gender and race can bias any conclusions you make about the individual involved.<sup>35</sup> As an extreme case of overgeneralization, if you have been bitten by a dog, you may develop a fear towards *all* dogs, which would be called a phobia. Such fear is quite irrational in the sense that it is very unlikely that the other dogs would bite you. This is a very concrete example of how thinking in terms of categories leads to suffering, as will be discussed in more detail in later chapters.

There are actually good computational reasons why overgeneralization occurs. Learning only a limited number of categories and using them without too much reserve means that knowledge gleaned from all the instances of each category can be pooled together; at the same time, the computational load is decreased. If you actually had enough data from all the dogs in the world, and had unlimited computational capacities, you would know some of them are safe while a few are not. However, data and computation are always limited, so some shortcuts may be necessary—even if they increase your suffering. This is another theme that we will return to over and over again in this book.

## The many faces of frustration: Summarizing the mechanisms of suffering

With this framework of information-processing in two systems, we can better summarize the previous chapters, in which we saw several computational ideas related to suffering. We started by considering two basic mechanisms: frustration, and threat to the person or the self. Later, we argued that threats to the self can be seen as special cases of frustration, namely frustration of self-needs (Chapter 6). Thus, we obtained a unified view in which suffering is based on error signals typically related to frustration of some kind.<sup>36</sup> We first defined frustration as not reaching a goal (Chapter 3) and later in terms of reward loss and reward prediction error (Chapter 5). In fact, these two kinds of frustration align well with the dual-process theory—slow vs. fast or GOFAI vs. neural networks—considered in this chapter.

However, there is much more than just two kinds of frustration. Reality is of course a bit more complex than such a clean division into planning and habit-based actions. Consider a case where you are yourself going to fetch the orange juice from the fridge. You formulate a plan which involves high-level actions such as going to the fridge, opening the door, etc. Once you are in front of the fridge, your habit-based system suggests you open the door by a certain sequence of muscle contractions which you have performed hundreds of times and

<sup>35</sup>(Freeman and Johnson, 2016). In that case, a further problem is that the categories may operate using stereotypes (which may not be factually accurate to begin with), which means that the generalization is even more wrong.

<sup>36</sup>It is also possible to see the connection between frustration and threats from the opposite angle: frustration could be seen as a special case of a threat to the self. I will tentatively sketch such a theory here. Perhaps failing in a task implies a threat to one's "self-image" that Cassell talks about, and which is related to the self-evaluation of Chapter 6. Frustration would imply that the agent's self-image is not correct, and that it has to change its self-image to something where it is less competent than it thought. Alternatively, any frustration could be considered to imply a threat to survival in the sense that it suggests that the agent's decision-making system is not optimal and could lead to serious problems in the future. In these ways, frustration might be reduced to a special case of the threat of intactness of the person, as in Cassell's definition of suffering. This would be in line with the thinking prevalent in Mahayana Buddhist schools, where self is seen as the source of all desires and all suffering. However, such a definition includes terms that lack a very precise definition in the framework of this book, in particular "threat", so further work is needed to formulate this approach in detail. (I'm grateful to Michael Gutmann for suggesting this interpretation).

which has become quite automated.

Now, suppose you follow the habit-based system and pull the door handle, but the door does not quite open. This kind of “frustrates” your habit of opening the door. But do you suffer? Probably not very much; you just pull again with more force, and if it opens, you hardly register anything out of the ordinary happened. In contrast, if you don’t get the juice at all—because the door is somehow broken and does not open at all— your long-range plan is frustrated, you will definitely suffer. And you really should suffer: All that planning and even the walking was in vain. A strong error signal has to be sent throughout your brain, and that is suffering.

### *Frustration on different time scales*

This example points out one important aspect of action selection: its temporally hierarchical nature, involving simultaneous computations on different time scales.<sup>37</sup> In the brain, there are also processes operating at many different time scales. So, some form of frustration can be operating on many different levels simultaneously. In one extreme, the agent may be planning long action sequences, and if they fail, frustration ensues in the sense of not reaching the goal. In the other extreme, a habit-based reinforcement learning system builds predictions on what kind of rewards or changes in state-values are associated with different actions, and computes whether there is reward loss or an RPE. Predictions are made on a millisecond time scale as well as on the time scale of days if not years. Each such time scale has its own learning mechanism using its own errors.<sup>38</sup>

Such division into time scales brings us to the concept of intention—defining intention as commitment to a goal, as discussed in Chapter 3. The point in intentions is to partly resolve conflicts between long-term and short-term optimization. I can have many desires simultaneously and spend some time thinking about each of them, and perhaps even planning each of them to some extent. But I’m not really hoping to reach *all* the goals related those desires. Once I decide to commit to one of the goals, that is what sets the goal which can then be frustrated. I would argue that in the case of planning, frustration is not so much due to desire itself but to the ensuing intention. This is in line with the more elaborate expositions of the Buddha’s philosophy on suffering which divide desire into initial desire and a later part called *attachment* (also translated as “clinging” or “grasping”). Attachment is a process where after an initial feeling of desire (“Nice, chocolate, I would like to have it”), you firmly attach to the object of your desire (“I must have that chocolate”). This distinction is, I think, similar to the distinction between desires and intentions: Buddhist philosophy suggests a central role for attachment, or intention, in the process which creates suffering. While such attachment or intention is not necessary for frustration to occur, I propose that it greatly amplifies it. This is logical because intentions consider longer time scales, and thus an error with intention is more serious, since more time and energy was lost in formulating and executing the plan that failed.

In fact, there is something that works on an even longer time scale: the frustration of self-needs treated in Chapter 6. Self-needs often work on time scales of days, months, even years. Casual observation suggests that they produce some of the very strongest frustrations. Different time scales may further be related to van Hooft’s different kinds of frustration discussed in Chapter 2: frustration of biological functioning, of desires and emotions (in his terminology, which may be different from this book), of more long-term life goals, and

<sup>37</sup>This is called hierarchical control (Poole and Mackworth, 2010), hierarchical planning or hierarchical task networks (Georgievski and Aiello, 2015; Nau et al., 2003) or hierarchical reinforcement learning (Sutton et al., 1999; Dietterich, 2000; Botvinick, 2012). A related paradigm is given by what is called “options” in reinforcement learning, see (Sutton and Barto, 2018, p. 461).

<sup>38</sup>(Hari et al., 2010; Botvinick, 2012). Using RPE instead of reward loss simplifies the situation to some extent, since RPE considers the total future reward and is thus less dependent on the definition of the time scale, as explained in footnote 20 in Chapter 5.

even of the sense of the meaning of one's existence.

### *Suffering based on desires, expectations, and general errors*

Another major difference between the two kinds of frustration—not reaching the goal and computation of a reward loss—is that reward loss is based on violation of *expectations*, while not reaching the goal is in line with the typical definition of frustration as not getting what one wants, i.e. violation of *desires*. One way of resolving this is to consider that the term “expectation” may have a slightly different meaning in the case of action selection, and especially planning. The agent is executing a plan in order to get to the goal state, and it is in that sense “expecting” to get to that goal state. Earlier, we saw (page 19) how Epictetus talks about desire “promising” the attainment of its object. Thus, the expectation related to planning could simply be defined as the goal state being reached.<sup>39</sup> Then, reward loss would be the same as the basic frustration of not reaching the goal, that is, the object of its desire (by the definition of Chapter 3).

Alternatively, we could see frustration (of desires) and reward loss (compared to expectation) as two distinct, if closely related phenomena, both of which produce suffering. What they have in common is that some kind of *error* occurred. On the computational level this is mainly an alternative way of defining the word “frustration”. But it opens up the possibility of a very general viewpoint where the connection between suffering and error signalling does not need to be concerned with goals or rewards at all. We all know that it is unpleasant if we expect something and then it does not happen, even if the event we were predicting was neutral in the sense of providing no reward. Thus, it is possible that there is some kind of suffering in almost *any* prediction error.<sup>40</sup> Most interestingly, it has been proposed that dopamine signals prediction errors for events not related to reinforcement, so it might provide a neural mechanism for general signalling of errors.<sup>41</sup>

In fact, I am tempted to think that desire (or aversion) in itself, especially when combined with intention, can immediately create some kind of suffering even *before* any frustration. Perhaps the internal representation of a goal state which is *different* from the current state is an error that automatically leads to the triggering of an error signal, and to suffering. (Alternatively, it could be that the system predicts that there will likely be frustration, and this leads to suffering by some kind of anticipation.) I will not develop this point any further here, but I point out that aversion in the form of irritation is clearly a kind of suffering in itself, while it is somewhat inadequately explained by the developments in earlier chapters.<sup>42</sup>

Frustration is further modified by the context. If you are deliberately engaged in the learning of, say, a new skill, errors are quite natural and you are likely to feel less frustration; in a sense, you are expecting that there are errors. Or, if your prediction of the reward is uncertain, i.e. only very approximate, the frustration is likely to be weaker. We will have much more to say about such effects in later chapters.

To recapitulate, we see quite a wide spectrum of frustration-related error signalling. Not reaching a goal, not getting an expected reward, or making an error in predicting any event, can all be seen in this same framework.

<sup>39</sup>The difficulty of defining expectation was earlier discussed in footnote 14 in Chapter 5.

<sup>40</sup>On the other hand, if we expect something unpleasant to happen, and it does not happen, we feel relief which is clearly not suffering; in general, obtaining more reward than expected may be the very definition of pleasure (see footnote 17 in Chapter 5). Perhaps, in that case, the unexpected positivity (of reward) overrides the inherent suffering in prediction error. For research on relief from pain, see Seymour et al. (2005); Leknes et al. (2008).

<sup>41</sup>Takahashi et al. (2017); Redgrave and Gurney (2006).

<sup>42</sup>The question of aversion will be treated later, in particular footnote 13 in Chapter 14. Also, Chapter 8 will propose that desire as well as some forms of aversion can be seen as “interrupts”, which may produce suffering by a special mechanism (page 95).

They work on different time scales, and use different systems in the dual-process framework. It seems that particularly strong suffering is obtained by frustration of planning, and even stronger by frustration of self-needs.

*Why there is frustration: Outline of the rest of this book*

We also need to understand why there is frustration in the first place. On some level, it is obvious that we cannot always reach our goals, or get what we want, if only because of the limitations in our physical skills and strength: We cannot move mountains. The world is also inherently uncertain and unpredictable, so even the perfect plan may fail because something unexpected happens. Yet, more interesting for our purposes are the cognitive limitations. As argued earlier, cognition is something that can be relatively easily intervened on, and modified to some extent. Thus it is more feasible to reduce suffering by focusing on the cognitive mechanisms, instead of trying to develop devices that physically move mountains. Therefore, it is crucial to understand in as much detail as possible how various processes of information-processing contribute to suffering.

We have already seen several information-processing limitations that can produce or amplify frustration. For example, planning is difficult due to the exponential explosion of the number of paths, which means our plans may be far from optimal. We need a lot of data for learning: data may be lacking to build a good model of the world, or to learn quantities such as state values. Categories are often used in action selection—in particular, if the world is divided into states—but these categories may not even be well-defined. The cognitive system may be insatiable and always want more and more rewards. There are several self-related needs which can create particularly strong suffering by mechanisms related to frustration.

Next, Part II goes into more depth regarding such limitations that produce suffering, focusing on the origins of uncontrollability and uncertainty. Later, Part III will consider methods for reducing suffering, mainly by reducing frustration. I will summarize all the different aspects of frustration in a single “equation”, and propose various methods, or interventions, to reduce frustration based on the theory of Parts I and II. Such methods will largely coincide with what Buddhist and Stoic philosophy propose, and include mindfulness meditation as an integral tool.

## **Part II**

# **Origins of suffering: uncontrollability and uncertainty**

**The second part will consider how uncontrollable the world as well as the cognitive system itself are, and how an agent's perceptions and thinking are uncertain and can even be called illusory.**



## Chapter 8

# Emotions and desires as interrupts

Part II of this book is about better understanding why there is suffering and what increases it. Since we saw earlier that suffering can fundamentally be seen as frustration, the question is what factors increase frustration. Part II analyses frustration in terms of uncontrollability and uncertainty (which is related to unpredictability). These properties make errors in action selection likely, and thus lead to frustration. Even the mind itself is seen as uncontrollable, since it has multiple processes operating at the same time, in particular emotions (Chapter 8) and wandering thoughts (Chapter 9). Further, perceptions are uncertain due to incomplete input data as well as a faulty prior model of the world (Chapter 10). The difficulties of communication between different brain areas or processors create a further loss of control (Chapter 11). Ultimately, we need to confront the problem of consciousness (Chapter 12) which creates a kind of a virtual reality where painful events are simulated again and again. For a sneak preview of what the system will look like in the end, the reader can have a look at Figure 13.1 on page 150.

\*\*\*\*\*

In this chapter, we start this investigation by looking at the concept of emotions. Anybody pressed to give sources of suffering would probably give a list of such phenomena as fear, disgust, sadness, and perhaps anger. Those are actually some of the most typical emotions in the terminology of neuroscience and psychology. If we are to understand suffering, we have to understand how such emotions are related to it. In Chapter 6, we already saw how fear is an essential part of self-related suffering, being due to threats to the self. Yet, fear is more than that: When assailed by fear, you forget everything else you were doing, you focus your attention exclusively on whatever caused your fear, you try to figure out how to get rid of it, and, eventually, run fast. These are examples of the aspects of emotions we investigate in this chapter.

We discuss how emotions can be seen as information processing and signalling, focusing on fear as a prime example. The main focus here is how emotions capture attention and interrupt ongoing processing. Another aspect is that emotions trigger basic, pre-programmed behavioural sequences, such as running away. Importantly, emotions are something that reduces any control we have on our minds and bodies, which is one leitmotiv in this part of the book. We also see that desires have similar interrupting qualities.

## Computation is one aspect of emotions

Some readers may wonder what emotions have to do with artificial intelligence. Surely, we can program an AI or a robot to function using purely “rational” procedures: Maximize expected reward and act accordingly, within the limits of the information the AI has, and as far it is computationally possible. Why would we need to introduce anything “emotional” in the system?

First, we need to understand what the word “emotion” means. Unfortunately, very different definitions are used, and there is no generally agreed definition even in the limited context of neuroscience and psychology.

### *Emotions have many components*

The most comprehensive definitions define an emotion as a complex of several different components. For example, if you feel fear, you will have a particular facial expression, you may scream, and your body will undergo physiological responses such as increased heartbeat. Next, your cognitive (i.e. computational) apparatus will start planning how to escape from the situation, and indeed, pre-programmed behavioural routines such as fleeing may be activated. While all this is happening, you will also feel afraid, in the strict sense that you have the conscious experience of being afraid.

As with almost any phenomenon in neuroscience and psychology, some emphasize the behavioural aspect of emotions, while others concentrate on more internal phenomena, including information-processing—usually called cognition in this context. Emotions are further characterized by a feeling tone: often negative (as in fear) but sometimes positive (as in joy). The feeling tone, technically called “valence”, is seen as the core of emotions by some, providing motivation for action. Yet others think that what defines an emotion is the conscious, subjective experience, such as feeling afraid.

In this book, I take an approach where all the aforementioned components together constitute an emotion.<sup>1</sup> Nevertheless, I focus on the computational, information-processing aspect of emotions, in line with the general approach of this book. Such information processing is often reflected in behaviour, and at least in humans, often leads to a subjective conscious experience, but I don’t explicitly consider behaviour or conscious experience in this chapter. The key question here is, how is information processed in what we call emotions; what is special in that information-processing when we feel, for example, fear or disgust?

### *Emotions help when computation and information are limited*

The starting point here is that emotions are needed because of the limited information available and the limited computational capacity. If an agent knew exactly everything that happens in the world and had unlimited computational power, perhaps it would not need emotions. A planning system would decide the best course of action—and it would really be the best course of action. However, in reality, things happen that we didn’t expect. It is because the agent does not know everything about the world (limited information), and the planning system cannot compute all the possible courses of action (limited computation). This is of course a narrative running through all AI and all neuroscience, but it is worth repeating.

---

<sup>1</sup>Such a multi-component definition of emotions, for example by Scherer (2009), is wide-spread in psychological literature. In contrast, in neuroscience and AI emotions are often defined more narrowly, or not properly defined at all. The list of such components given in the text is not at all complete, in particular bodily responses (Nummenmaa et al., 2014) and social aspects (Nummenmaa et al., 2012) could be added.

This chapter will show various ways in which emotions help in information-processing under such limitations. One thing that the limitations above imply is that some kind of monitoring of unexpected events is needed, as well as a system for changing plans as they are detected. This is the role of interrupts, which is one of the themes of this chapter, and one of the specific functions of emotions. Such interrupts often also trigger pre-programmed action sequences or plans that have been found useful by evolution, or the programmer, which is another aspect of how emotions help in steering an agent's behaviour.

## Emotions interrupt ongoing processing

Suppose you (or a robot) are walking home on a street you know. While walking, you may be planning what you will be eating tonight (the robot might be just concentrating on the walking because that's difficult enough for it). Now, a car suddenly appears and comes fast in your direction. What you need to do to survive is two things. First, your perceptual system has to detect that something unexpected and potentially dangerous is happening. Second, the fact that something potentially dangerous is happening must be broadcast to the whole system; you have to stop thinking about what you will eat, and you have to stop following the route back home. Thus, you interrupt all ongoing activities, including your current train of thought. Instead, you have to use all your cognitive resources to figure out what to do, how to jump to safety and when.

The important new twist here is that once the sensory systems realize something suspicious is happening — even if they don't exactly know what — they have to send a kind of an *alarm* signal to other parts of the brain. In particular, the system responsible for executing action plans must be interrupted; in computer science, such a signal from one process to another is typically called an interrupt. These functionalities go much beyond the mere “cool” perception that a car is visible and coming in your direction.

A separate alerting mechanism with the capacity to stop ongoing activities and reorient computation is the core of the *interrupt theory* of emotions originally proposed by Herbert Simon in the 1960s.<sup>2</sup> The key idea in this theory is that being an interrupt is what distinguishes an emotion from ordinary information-processing. The interrupt theory explains why emotions have particularly powerful attention-grabbing properties; that is the whole point of emotions according to this theory.<sup>3</sup> Such an interrupt system is particularly important since earlier in Chapter 3 we argued, following the belief-desire-intention theory, that an agent needs to commit to a single plan instead of jumping from one plan to another. Commitment is useful, but it should not be blind: interrupting a plan must be possible.<sup>4</sup>

### *Pain, disgust, and fear*

At the most elementary level of interrupts, we actually find simple physical pain. Although we don't categorize it as emotion, pain is clearly a signal or a process that has such an interrupting quality. It is broadcast to the

---

<sup>2</sup>(Simon, 1967; Oatley and Johnson-Laird, 1987)

<sup>3</sup>The concept of attention is mainly elaborated in later chapters, but anticipating them, I need to point out that interrupts are closely related to a specific kind of attention which is *bottom-up* attention, see footnote 16 in Chapter 10. I don't elaborate the connection between attention and interrupts here, and I use the word “attention” casually, in its everyday meaning.

<sup>4</sup>The attention-grabbing properties of emotions are well understood by the designers of social media platforms. The more the news and updates evoke fear or anger, the more attention the user pays to the platform. The avowed primary goal of some such platforms is engagement, which is basically one aspect of attention. Some negative side-effects of designing such systems should be well-known to everybody by now.

whole information-processing system; all ongoing behaviours are typically suppressed, and the organism uses most of its resources to get rid of the cause of the pain. Pain is, in fact, the most fundamental, as well as the strongest kind of interrupt. It has to be so, because it is the signal which is the most relevant for the intactness and even the very survival of the organism. The alert is about a physical, chemical, or biological danger to the organism—tissue damage in the terminology of Chapter 2—that typically comes from outside.<sup>5</sup> It requires urgent action, such as withdrawal away from the object that caused the pain. Reflexes like this are present even in very simple organisms, and should be programmed even in reasonably simple robots. You don't want an expensive robot to break down the very first day because it doesn't understand what kind of actions are dangerous to itself.

Disgust is conventionally classified as an emotion, although it is closely related to pain. Disgust is triggered by perception of substances which are likely to be toxic or transmit diseases. Again, current processing is interrupted to direct attention to that substance and how to avoid it. Disgust is often a very primitive emotion: for example, disgust at the smell of rotten food is very close to physical pain. This is natural since disgust is about protecting the organism from something not very different from tissue damage. However, disgust has also more abstract forms in the case of disgust at morally condemnable behaviours.<sup>6</sup>

More complex organisms are able to *predict* impending danger at a much greater distance and temporal delay, as discussed at the end of Chapter 6. While disgust, and even pain, already have such a predictive quality in primitive form, complex organisms can predict risk of damage before the pain or disgust systems are activated. The signal related to such anticipated danger is *fear*, which interrupts ongoing activity and directs processing to avoidance of the dangerous object. This interrupting viewpoint on fear is different from our discussion of fear in Chapter 6, where we linked fear directly to suffering through the concept of self. Thus we see explicitly how emotions have many components or aspects even regarding computation.

## Desire as an emotion and interrupt

Interrupts can also be useful when there is no danger visible, but rather an opportunity to obtain some kind of reward. Casual observation tells us that something very similar to an interrupt happens when you see an object that you really like and want. You are assailed by an acute, “burning” form of desire. While in neuroscience and psychology desire is usually not considered an emotion, there has always been some doubt on whether such a distinction is justified. Acute, burning desire actually squarely sits in the domain of emotions as far as the interrupt theory is concerned.<sup>7</sup>

In chapter 3, we first defined the desire system as something that suggests goals to the planning system. But we didn't go into details on how the desire system actually works: How can it identify states which are easy to reach while having a high state-value? I think the whole point in the computations related to desire is that they happen as a dual process. When desires suggest goals for a planning system, they have to do it based on fast neural network computations in order to usefully complement planning. As we saw in Chapter 7, neural networks, such as those in AlphaGo, can be trained to output approximate solutions to the computations of state-values and similar quantities needed for planning. It is likely that the computations underlying desire

<sup>5</sup>In line with the interrupt theory, Craig (2003) argues that pain should be seen as an “emotion” as it includes “a behavioral drive with reflexive autonomic adjustments” unlike plain sensory processing.

<sup>6</sup>(Chapman and Anderson, 2012)

<sup>7</sup>(Oatley and Johnson-Laird, 1990)

are based on such neural networks, which suggest candidate states that are likely to be easily accessible while having a high state-value.

### *Elaborated-intrusion theory*

A psychological theory that is very compatible with such goal-suggestions by neural networks, and combines it with an interrupting quality, is the *elaborated intrusion* theory of desires. As its name implies, it considers desire as a computational process that *intrudes* your mind: it invades your information-processing system so that you lose control, at least initially. You are not able to think about anything else and keep planning courses of action regarding that object of your desire. Such ensuing compulsive planning is the *elaboration* part of desire.<sup>8</sup>

Everybody has experienced such intrusions. You see a sexually attractive person, and you cannot think about anything else for a while. Or, you see your very favourite brand of chocolate in a supermarket, and you can hardly resist taking it in your hand and putting it into your shopping basket. You may be devising all kinds of sophisticated plans to get the object of your desire, forgetting completely what you were actually supposed to be doing. Thus, at least in humans, the simple neural networks computing desire can be in conflict with deliberative planning processes. This emphasizes that desire can take control of the mind, inexorably turning our attention towards the object of the desire. With such really “hot” desire, which could be called “irrational” and strongly affective, there can be a conflict between “reason and passion”—which is perhaps a poetic expression for the dual-process character of the information-processing system.<sup>9</sup>

### *Valence*

Such a dual-process approach above brings us close to another interesting concept: valence. In psychology, valence is a technical term describing the intrinsic positive-negative, pleasure-displeasure, or good-bad axis of states or objects. From the viewpoint of subjective human experience, valence means whether feelings are positive or negative: positive valence is associated with pleasure, negative valence with displeasure. Valence is closely related to liking: we could equate liking and valence, saying that we like things which have a positive valence and dislike things which have a negative valence. Alternatively, valence can be defined based

---

<sup>8</sup>I follow here Kavanagh et al. (2005). A closely related model which talks about “impulses” instead of desires, and explicitly links them to a dual-process theory, is presented by Hofmann et al. (2009). Similar ideas can be found in consumer research; Belk et al. (2003) in particular contrast desires and what they call “needs” as: “We burn and are aflame with desire; (...) we are tortured, tormented, and racked by desire; (...) our desire is fierce, hot, intense, passionate, incandescent, and irresistible; (...) Needs are anticipated, controlled, denied, postponed, prioritized, planned for, addressed, satisfied, fulfilled, and gratified through logical instrumental processes. Desires, on the other hand, are overpowering; something we give in to; something that takes control of us and totally dominates our thoughts, feelings, and actions.”

<sup>9</sup>To clarify and recapitulate: We can define desire in different ways on the hot-cold axis. In the coldest definition, desire is simply a preference for some states, essentially just another way of saying that some states are rewarding or have higher state-values. You might say, for example, that you want to see Kyoto one day, but saying that does not necessarily arouse any feelings, and launch any deliberations in your brain. A slightly less cold definition says that desires propose goal states for a planning system, thus possibly launching computations to attain such a state. The definition in the elaborated-intrusion theory is quite hot, emphasizing the interruptive quality of those computations. An even hotter definition, not pursued here in detail, might further add a subjective, conscious experience of burning with desire, but this is outside of the computational modelling framework we take here, and presumably only applicable to humans and higher animals. — To emphasize the difference between different kinds of desire, such a “hot”, compelling desire is sometimes called *occurrent* desire, while the kind of cold, long-term rational desire that simply expresses a preference is called *standing* desire. I prefer to talk about “interrupting” desire instead of *occurrent*.

on behaviour: humans as well as animals approach and try to obtain states which have positive valence, and avoid things and states with negative valence.<sup>10</sup> Desire is thus usually directed towards states that have positive valence.<sup>11</sup>

In our framework, valence is closely related to the quick evaluation of any state or object by the neural network that computes approximations of state-values. (I shall not attempt to give an exact definition of valence or liking here since it is not particularly important in what follows.) When you see chocolate, its high valence is reflected in your neural networks predicting high state-value if you reach the state of eating it. Thus, valence computations are necessary for interrupts based on desire. In Chapter 13 we shall discuss how the sequence valence-desire-intention is important in Buddhist philosophy: just like in the present discussion, it is valence that leads to desire, and further to intentions and frustration. In that sense, the valence computations are the very root of suffering.

## Emotions include hard-wired action sequences

One reason for having interrupts is that they often launch “hard-wired” programs, or sequences of actions for specific situations. Many emotions are characterized by their specific, relatively rigid programs.<sup>12</sup> The action sequence is, in fact, the aspect that most visibly distinguishes which emotion is taking place. In the case of fear, the typical action is to choose either freezing or fleeing. Disgust leads to immediate rejection and avoidance of the substance triggering the emotion. In animals, such programs are evolutionarily quite old: humans have largely the same action programs as dogs.<sup>13</sup>

The point is that some simple action sequences are particularly useful and universal, so it is a good idea to have them readily stored in the system so they can be executed quickly, without any need for elaboration. This is in stark contrast to the main processing being interrupted, which is often a result of long elaboration. In fact, plans may take quite a while to formulate, which makes planning less useful in an emergency situation. Furthermore, such emotion-specific action sequences may be very difficult to learn. For example, anything related to self-preservation is difficult to learn by reinforcement learning, since when the agent realizes that the current situation is lethal, it is too late. Therefore, it is important to have them readily programmed in the system—meaning genetically transmitted in humans.

Anger is another fundamental example of an emotion which clearly has its own hard-wired action sequence. It also has a particularly strong social quality: real anger in the sense of interrupt is usually associated with other people. While you might say that you are angry about bad weather, that is not much more than

---

<sup>10</sup>(Colombetti, 2005)

<sup>11</sup>Usually, people want things that they like, and vice versa. However, recent research has found that in some cases, people can want things which they don't like (Berridge and Kringelbach, 2015)—in the precise sense that those things do not produce physiological pleasure reactions. This phenomenon is one of the underlying mechanisms in drug addictions: An addict may want and consume the drug without actually deriving any pleasure from it. In fact, such desires don't even need to be conscious in humans. See also footnote 24 in Chapter 5.

<sup>12</sup>Simon's interrupt theory was elaborated by Oatley and Johnson-Laird (1987) by proposing how different emotions correspond to different action sequences. Going further in that direction, we find Frijda's theory of emotions as “action readiness”, meaning the preparation for movement or action (Frijda, 2016). Frijda's theory sees this as the main distinguishing feature of emotions, instead of their interrupting character. From the computational viewpoint, it could be argued that any simple neural-network-based reinforcement learning agent can trigger such action readiness, and it is difficult to see what would be special about emotions if they were defined as simply action readiness.

<sup>13</sup>(Gross and Canteras, 2012)

ordinary frustration. We shall not consider anger in any detail here because such social aspects are completely beyond the scope of this book and would require more complicated theory, in particular game theory. Let me just mention the basic idea, which is that anger is a special hard-wired action sequence that protects the agent from attacks by creating a credible threat of a robust retaliation that would inevitably be triggered the case of being attacked.<sup>14</sup>

It is now useful to contrast emotions to habits, in the wide sense used in Chapter 5. Habits are often triggered by some environmental stimuli—a bit like interrupts—and lead to a fixed kind of behaviour—a bit like the rigid action sequences we just mentioned. In these two ways, habits have some similarities to emotions. However, habits are not really interrupts. Perhaps when you walk on the street you have the habit of humming a tune to yourself. However, it rarely happens that you stop whatever you’re doing because you suddenly feel an irresistible urge to start humming. Habits don’t have the power to capture your attention and interrupt current plans.

## How interrupts increase suffering

We have seen that a number of phenomena, which are often considered separate in psychology and neuroscience, share the important characteristic of being interrupts. Pain, emotions, and desire can all be seen in this computational framework.<sup>15</sup> But many emotions include a lot of suffering. If such emotions are essentially just interrupts, why would there be so much suffering involved?

I propose there is a reason why interrupts create suffering directly, by themselves: the interrupting system uses the pain signalling system. In fact, most emotions discussed here are negative, they hurt, and this suggests they must use the pain system, like suffering (“mental pain”) in general.<sup>16</sup> Making the body feel pain is an evolutionarily primitive way of grabbing the attention of the whole cognitive system, as was discussed in

---

<sup>14</sup>Imagine a gangster comes to you and asks you to give him all your money. The rational thing to do would be to give the money. This is rational in the sense that otherwise, he might kill you or inflict some bodily harm, and certainly it is better for you to just give the money. However, this behaviour has the downside that then the gangster can come to you any time he wishes and always take your money. An evolutionary explanation of anger is that it is a program that makes you behave irrationally. In this case, you would just get mad, and physically attack the gangster, even if you know he will kill you as a consequence. Surprisingly, having such a program may be good from an evolutionary viewpoint, because if the gangster knows you have such a program installed, he might decide not to bother you. It is not good, from an evolutionary viewpoint, to actually attack the gangster; what is good here evolutionarily is having such a program installed, and signalling this to the gangster. If the gangster knows about the program, it may never be actually used, because it works as a powerful deterrent. This is a well-known game-theoretic model in evolutionary theory, it was originally used for modelling the behaviour of animals who fight over mating opportunities, territory, or other scarce resources (Smith and Price, 1973; Pinker, 1999), which is why it is often called the hawk-dove game (Hirshleifer, 1987; Nowak et al., 2016). It is actually equivalent to another game-theoretical model called the game of chicken, which, despite sharing an avian name, has a very different story and motivation behind it. — Let me also note that social interaction creates many further emotions, such as shame and guilt, some of which are moral emotions (Haidt, 2003); that is, they enforce behaviour conforming to ethical norms.

<sup>15</sup>While considering pain, basic emotions, and desires in a single framework is not usual in neuroscience, in recent neuroscience literature, ideas similar to the interrupt theory use the distinction between a planning system (“Model-based reinforcement learning”) and a fast system with automated reactions (“Model-free reinforcement learning”). If a system has such two systems, like the brain, a crucial question is how to divide tasks between the two systems, i.e. which one to use to respond to any particular situation (Daw et al., 2005). If fast action is required, you obviously need to use the fast system, and if there is no hurry, you can spend some time in planning, but choosing which to do is a complicated problem. The process leading to the decision to use the fast system is then not very different from the mechanisms postulated in the interrupt theory (Bach and Dayan, 2017). However, the starting point of the interrupt theory is to answer the deeper question of *why* it is useful to have two such systems in the first place.

<sup>16</sup>(Papini et al., 2015; Eisenberger and Lieberman, 2004)

Chapter 2. Interrupts need, by their very definition, to achieve such an attention-grabbing effect, so using the pain system is even more natural than in the case of, say, frustration. In fact, the signal that triggers an interrupt can be interpreted as a special kind of *error signal*, and thus it fits in our general framework of suffering based on error signalling.<sup>17</sup> (Positive emotions are a rather different story, and not considered here.<sup>18</sup>)

Another important problem with interrupts is that they *reduce control*, which is one of the main themes in the following chapters. A crucial part of the interrupt theory is the idea that the interrupts are automatic and largely irresistible. For example, many people would be so much happier if they could just consciously decide to switch off their fear system. But the point is that interrupts are outside of conscious control: They have to be so, because very often they need to interrupt conscious thinking and consciously controlled action. If you could somehow weaken interrupts so that they don't disturb you, they would be useless: it would be like switching off a fire alarm system because it is too loud. In a scary situation, fear will appear together with its inherent suffering, no matter how much you try to control it. We saw already in Chapter 7 how the dual-system structure of the brain means that the fast, unconscious fear system usually prevails. (We will return to the question of conscious control, or lack thereof, in Chapters 9 and 11.) The same happens with desires: The fast computations of valence and values by neural networks will "intrude" and interrupt other processing, directing all the processing towards the object of the desire. Such interrupts are even more annoying if they interrupt activity that would have created pleasure, for example when you are in a "flow", fully engaged in a rewarding and meaningful activity. In this sense, interrupts greatly *increase* suffering, by increasing desires, aversion, planning, and frustration.

Such reduction of control might not be a bad thing if the interrupts were somehow optimally tuned to reduce suffering. However, another problem with the interrupt system is that its design parameters are often questionable from the viewpoint of suffering. To begin with, the system that triggers interrupts does not care about our subjective suffering, only about our evolutionary fitness. Evolution makes us consider harmless things as dangerous, worth triggering an interrupt, if they are threats to our evolutionary success. Sexual jealousy, and the ensuing rage, is one example, where (from a male perspective) the evolutionary "danger" is that one might end up raising a child who is not one's own and does not spread one's genes. Yet, that is hardly a problem from a contemporary viewpoint: it is in fact very common in modern families.

What's more, the system may not actually be very good at maximizing evolutionary fitness either. As we saw earlier, evolutionarily developed neural mechanisms may not be well adjusted to our current society, since they may come from the legendary "African savannah". In the case of fear, for example, we tend to be afraid of snakes or spiders, but not so much of cars, although cars are much more dangerous at least in modern cities.

---

<sup>17</sup>Thus, such negative emotions seem to introduce a mechanism for suffering which may be a bit different from what has been discussed previously, at least in the sense of introducing a new kind of an error signal. But it is also possible that it is not really different from the frustration-based and self-based mechanisms treated in the preceding chapters. One can argue that just as we already reduced fear to frustration in Chapter 6, we can reduce the other negative emotions to *frustration of self-needs*. If an interrupt is triggered, that means something went "wrong" from the viewpoint of survival or threat-avoidance, and thus such self-needs were frustrated. What matters really from the viewpoint of this book is which theoretical approach is useful from the viewpoint of interventions. Frustration is an error signal on which we can actually intervene, and many methods will be discussed in Chapters 14 and 15. It may be more difficult to intervene on the interrupts of fear and disgust, for example (but Chapter 15 discusses on page 191, some possibilities related to threat and on page 184 related to desires). So, it might be more practical if the framework of frustration (including frustration of self-needs) were sufficient to explain the appearance of suffering even in the case of emotions. I leave details of such connections for future work to elaborate.

<sup>18</sup>Let me just mention that Fredrickson (2001) proposes positive emotions serve the role of enabling exploration of different action possibilities when the circumstances are safe and not even remotely life-threatening, thus leading to enhanced creativity and learning.



Another problem with the interrupt system is that if the interrupts are excessive and disrupt the normal function of the system too much, it may simply worsen the situation by making it more difficult to respond to the situation. Such problems are related to the fact that emotions and desires are short-sighted—as has been acknowledged by philosophers since antiquity—and may interrupt useful plans in a way that produces frustration because the interrupts fail to understand the long-term utility of following the plan. For example, an important function of pain is to attract the attention of the agent to the source of the pain, but if the person can think of nothing else than the pain, as often happens in the case of overwhelming fear or depression, he will not be able to find a solution to the situation. Or, if you are easily scared and are constantly interrupted by, say, harmless bugs, your performance in a meaningful pursuit may be hampered even though there was never any real danger to avoid.

### *Alarm systems cannot be universally optimal*

These questions are related to the general theory of designing alarm systems, which is considered in the mathematical theory called signal detection theory.<sup>19</sup> It is based on maximizing the expected payoffs, where payoffs are similar to rewards, describing how good (positive) or bad (negative) the results of a given action are. For an alarm system such as interrupts, there are two possible actions: trigger an alarm, or do not. The theory is related to the AI theory outlined in previous chapters, but with a different emphasis. An important lesson in this theory is that *there is no such thing as a universally optimal alarm system*. That is because the payoffs are different for different people, and different in different contexts, and may change over time.

Consider designing an alarm system yourself, in the form of a burglar alarm. You might start the design of a burglar alarm system by assigning a high payoff to detecting burglars—which sounds reasonable and innocuous. However, this means the system will not mind making false alarms, since you only give a strong payoff (reward) for detecting burglars, but you do not give any punishment for false alarms. To maximize reward, the system rationally decides to trigger an alarm if there is any hint of a burglar present. Eventually, your burglar alarm will constantly wake up everybody in the middle of the night. Realizing your mistake, you next give a really high reward for not giving false alarms. The result is that the system never gives an alarm because that's the perfect way to avoid false alarms, which are now strongly punished. In this case, the alarm system ends up being completely useless since it does not do anything. It is very difficult to say what the right compromise is: the alarm system should be sensitive but not too sensitive, and the right parameters are quite subjective and depend on the context. Evolution has programmed certain sensitivity levels in our interrupt system, but in light of this signal detection theory, it is not actually clear how optimal they were even for all our ancestors on the African savannah, let alone for modern city-dwellers.<sup>20</sup>

<sup>19</sup>(Green and Swets, 1988). This is a special case of statistical decision theory, typically considering the case of two possible options given by “present” or “absent” (regarding a threat), and focusing on the question of finding the right balance between false positives and false negatives.

<sup>20</sup>The sensitivity levels, or thresholds, can also be modified by experience to some extent. For example, if as a child, you saw something that made you really scared, you may lower the threshold a lot—commonly known as a phobia. This leads to another problem, widely recognized in clinical psychology: the payoffs change during an individual's lifetime as well. In adults, they may be very different from what they were in our childhood environment, while any learning of the payoffs may mainly happen as a child. The best survival strategy for a child in an adverse environment may be to be constantly afraid of other people; this may not be optimal anymore when the child grows up, and is in fact a possible source of psychiatric problems.

## Emotions are boundedly rational

Often emotions are contrasted with rationality and “cool-headed” decision-making; it is typically assumed that the best decisions are made when emotions are not at play. However, the viewpoints on emotions explained in this chapter show that emotions contribute to optimal decision-making and action selection. Emotions are useful from a rational viewpoint as soon as there are certain information-processing constraints; for example, if the planning system does not have time to consider all possible paths in the search tree. This is certainly true in any sufficiently complex AI system or animal. The viewpoint which considers emotions as necessarily irrational is in fact largely rejected in modern research.<sup>21</sup>

I have casually used the word “rational” here as well as in earlier chapters, but we need to think a bit more about what it actually means. Often, a decision is called rational if it is optimal in maximizing reward or a similar quantity given the information available to the agent. In other words, the decision of the agent (such as choosing an action) is the same as that made by an ideal, hypothetical agent with perfect information-processing capacities and the same information about the world as the agent in question. So, even a perfectly rational agent is not expected, in this definition, to make the very best possible decision, but the best possible given the limited information it has at its disposal. However, in reality, the information-processing power of the agent is limited as well, as we have indeed seen in many chapters of this book.<sup>22</sup> The case where information-processing power is limited as well leads to the concept of *bounded rationality*, also called computational rationality. It refers to decisions which are optimal given limitations in both the information *and* the computation available to the agent.<sup>23</sup>

Emotions, seen as interrupts or as automated action sequences, can largely be considered to strive towards bounded rationality. In both cases, emotions are information-processing routines or shortcuts which help in achieving as good outcomes as possible, given the computational restrictions and the limited information available. It is in this precise sense that we can say that emotions help in rational decision-making, and it is not justified to oppose rationality and emotions.<sup>24</sup>

Yet, emotions also have qualities that are in contrast to our everyday notion of rationality. In particular, they are not under conscious control. In this sense, they are similar to the neural networks in dual-process theories, and indeed we saw that connection above in the case of desire. The question of control is actually crucial from the viewpoint of suffering, as we will see many times in the following chapters.

---

<sup>21</sup> (Damasio, 1994; Scherer, 2011)

<sup>22</sup> Curiously, largely due to historical reasons, limitations in information available to the agent were always admitted, but computation was not supposed to be an issue in earlier work on rationality. I’m here referring to the classic work on statistical and economic decision theory in the first half of the 20th century, arguably culminating in the work by Von Neumann and Morgenstern (1944).

<sup>23</sup> (Simon, 1972; Russell, 1997; Gershman et al., 2015; Lieder and Griffiths, 2020)

<sup>24</sup> Another rather different information-processing function of emotions has been proposed as the “somatic marker hypothesis” by Antonio Damasio (Damasio, 1994; Bechara and Damasio, 2005). Somatic markers are defined as bodily responses to situations, learned from past experiences. If a certain situation has led to a bad outcome, you learn to associate such a situation with a bad “gut” feeling in your body. The somatic marker hypothesis thus shows how such feelings (here considered the essential part of emotions) can be used to improve planning by using them as heuristics. As we saw earlier, the central problem in action selection is the huge, exponential number of plans to consider. Using somatic markers as heuristics, you may be able to reject many of them based on such negative feelings and focus your search on the set of plans associated with positive gut feelings. Importantly, such “gut feelings” are generated by a very fast computation in a simple feed-forward neural network, thus speeding up decision-making and planning—not unlike the computations we linked to desire, valence, and dual-process action selection earlier in this chapter and Chapter 7.

## Chapter 9

# Thoughts wandering by default

The moment you lie down on a sofa to relax, your head starts developing different fantasies and daydreams, perhaps wondering why you did such a stupid thing yesterday, or planning what you want to eat tonight. Even when you try to meditate and not think about anything (which is a typical instruction for beginning meditators), you will almost inevitably find yourself thinking about something else after a while. There is a good reason why the human mind is often compared to a monkey in meditation traditions. It jumps here and there, making all kinds of noises, and never seems to rest. Likewise, based on his own method of introspection, David Hume concluded: “One thought chases another, and draws after it a third, by which it is expelled in its turn.”<sup>1</sup>

Thoughts that come to your mind when you are trying to concentrate on something else are called “wandering thoughts”. They have some similarities with emotional interrupts: they stop ongoing mental activity and capture attention. Thus, they reduce the control you have over your mind and, eventually, increase suffering. However, the computational underpinnings are quite different in the two cases. In this chapter, I discuss how wandering thoughts are related to the need to repeat experiences for the purposes of iterative learning algorithms, as well as planning the future through search in a tree. Thus, there is an evolutionary reason why we have wandering thoughts: they are not just pointless activity triggered by mistake, as it were.

### Wandering thoughts and the default-mode network

Wandering thoughts tend to appear whenever a person tries to focus on a single task or object for a long time. Everybody has encountered a situation, perhaps at school or at work, where she tries to concentrate on something but soon finds herself thinking about what she should say in a job interview tomorrow, or what she did on a previous vacation. Typical tasks where such *sustained attention* is necessary, but difficult to achieve, are driving a car on a highway, trying to read a book for an exam, or monitoring a screen as in air traffic control or surveillance. Importantly to the theme of this book, sustained attention is essential in most meditation practices. If you are lying on a sofa and have nothing else to do, meandering thoughts about various things are fine and sometimes even enjoyable. However, when you are actually trying to concentrate on a task, unwanted wandering thoughts reduce your performance of the task at hand.<sup>2</sup>

---

<sup>1</sup> (Hume, 1739), Section 1.4.6

<sup>2</sup> In the case where you have no particular task to perform, the spontaneously appearing thoughts may not be properly called “wandering”, but, for example, “spontaneous”. Some authors strictly reserve the word “wandering” for the case where the thoughts are intrusive in the sense of occurring against your will while you are trying to concentrate on some task, such as thinking about some un-

Psychological experiments confirm the ubiquity of wandering thoughts. Various experiments can be devised where the participant's task is to monitor a stream of information and report when a rare prespecified event occurs. In a typical experiment, you would be shown random digits (0 to 9), and you have to press a button when you see a target digit, say, 3. The experiment is deliberately designed to be boring so that the participant's mind will certainly start wandering at times. The basic idea of monitoring for an event that is rare is reminiscent of some of the typical real-life tasks listed above (e.g. driving a car on a quiet highway), where nothing much happens most of the time, and sustained attention is difficult. The experimenters would then use a method called experience sampling, which means they ask, at random intervals, whether the participant was focused on the task or whether they had wandering thoughts. It typically is found that the participant's performance on the task fluctuates between better and worse; this fluctuation largely reflects whether they had wandering thoughts at that particular time point or not.<sup>3</sup>

Such experiments can be conducted even when the participants are living ordinary everyday lives. The participants would have a device, such as a mobile phone, which asks at random intervals whether they were focused on whatever task you were performing (such as working, studying, cleaning, driving, etc.) or whether they had wandering thoughts (such as daydreaming, fantasies). It is typically found that during everyday life, the mind is wandering quite a lot: one third, or perhaps even one half of the time.<sup>4</sup>

### *Much of brain activity is spontaneous*

At the same time, modern neuroimaging confirms the prominence of various kinds of spontaneous brain activity, i.e. activity that “just happens” without any external stimulation or task being performed. In fact, an amazing finding in recent neuroscience is that if you measure human brain activity when the participants of the experiment are simply told to sit or lie still and think about nothing in particular, their brains are far from quiet. Technically, neuroscientists talk about “resting-state” to characterize such a state of not doing anything in particular, since the participant may think she is having a rest—but the brain is definitely not.<sup>5</sup>

A particular network in the brain is actually even more active during rest than during active tasks. It is called the *default-mode network* because it seems to be activated “by default”, i.e. when there is no particular reason for anything else to be activated.<sup>6</sup> It is also deactivated once the person is stimulated by, for example, sights or sounds from the external world, so that the brain actively starts processing incoming information.

The discovery of the default-mode network around the year 2000 was something of a revolution in human

---

related event tomorrow when trying to concentrate on reading a textbook. In this book, I use the term wandering thoughts a bit more liberally, sometimes including thinking that jumps from one topic to another when there is no particular task on which it is supposed to concentrate—as in lying on the sofa after work—since in real life, it is often difficult to draw the line between wandering and other spontaneous thinking.

<sup>3</sup>(Christoff et al., 2009)

<sup>4</sup>(Kane et al., 2007; Killingsworth and Gilbert, 2010)

<sup>5</sup>In terms of neural network theory, resting-state activity is enabled by the brain having *intrinsic dynamics* based on *recurrent* connections. Recurrent connections mean that the neurons are not arranged in successive layers where the signal just goes in one direction: Instead, the outputs of some neurons are fed back to other neurons that actually provided input to those neurons in question. The output can also be fed back to the outputting neuron itself. With such feedback, neurons can learn to sustain each other's activity: Neuron A activates neuron B, which by recurrency again activates neuron A, and so on. Even a single neuron can sustain its activity by sending feedback activation to itself (Hopfield, 1982; Hochreiter and Schmidhuber, 1997). Such recurrent connections are extremely common in the brain, while the most commonly used neural network models have no recurrent connections; this is an important discrepancy.

<sup>6</sup>For recent reviews, see Buckner et al. (2008); Raichle (2015), for the original articles, Shulman et al. (1997); Raichle et al. (2001).

neuroscience. It was completely at odds with the classical way neuroscience experiments were done: the experimenter would instruct an experimental subject to observe some stimuli (e.g. a sequence of digits as we saw above) and possibly perform a task at the same time (e.g. press a button when a target digit appears). Here, in contrast, you don't tell the participants to do anything and don't give them any sensory stimulation, such as showing pictures. Then, it is the default-mode network that becomes activated. In fact, since it is silenced (i.e. deactivated) by sensory stimulation and tasks, the experimenters had better not give any stimuli or tasks to be able to observe it.<sup>7</sup>

It is widely assumed that the default-mode network supports wandering thoughts.<sup>8</sup> That would explain why it is particularly activated when the subjects do not receive any stimulation and have no particular task: then, the mind will easily start wandering. It is likely that the default-mode network has other functions as well, although we don't know very well what they might be.<sup>9</sup>

## Wandering thoughts as replay and planning

The existence of wandering thoughts may feel completely normal to us, but actually, it is rather surprising that the whole phenomenon exists. Why should it be difficult to concentrate on one thing for a long time? Why cannot I just decide to focus on reading a textbook for an exam, say for two hours, without any interruption by any unrelated thinking?<sup>10</sup>

One intuitively appealing explanation would be that your active neurons—in the exam-reading example, those needed for reading—get “tired”, i.e., somehow run out of energy. Then, other neurons which are full of energy will be able to somehow steal the attention. While there may be some truth in such an explanation, it is not very plausible because sometimes you *can* concentrate without any problems on a task, especially on a task which is really engaging, such as reading a book you really like (not for an exam), or playing a video game. Furthermore, should not such fatigue of neurons rather lead to having no thoughts at all? It is more plausible that wandering thoughts are actually doing some useful computation—and that they are something that you would like to program in an AI.

So, let's think about what kind of computational problems could be solved by wandering thoughts. One problem we have seen earlier is that learning typically needs many repetitions of the inputs and the desired outputs, being based on iterative algorithms, as we saw in Chapter 4. Even the very same inputs and outputs may need to be presented many times to the learning algorithm. This is why one thing that modern AI systems have in common is that most of the computing capacity is used for learning. At the same time, planning takes a lot of time as well, as we saw in Chapter 3.

---

<sup>7</sup>Actually, it has further been found that many of the same brain networks that are intermittently active in various neuroscience experiments are also intermittently active in resting-state. Thus, the default-mode network is not the only network activated in resting-state, but the default-mode network is perhaps the only one that is *more* active in resting-state than in *any* kind of stimulation or task. When we talk about a “network” here, we mean more precisely that the activity in the same sets of voxels (and, presumably, the same sets of neurons) seem to be fluctuating synchronously (Damoiseaux et al., 2006; Fox and Raichle, 2007). Typically the analysis is made using a machine learning method called independent component analysis (Hyvärinen et al., 2001). Possibly the first study to provide such a decomposition to several networks was in fact based on data from anaesthetized human children, whose brains were scanned for clinical purposes (Kiviniemi et al., 2003).

<sup>8</sup>(Christoff et al., 2009; Andrews-Hanna, 2012)

<sup>9</sup>(Raichle, 2015)

<sup>10</sup>(van Vugt et al., 2015)

So, as much of the computing capacity as possible should be directed to these learning and planning activities. In particular, when the agent does not receive any special stimulation from outside, there is nothing important for it to do, and no urgent threats are detected, the computing capacity of the agent is free to be used for any internal processing based on previously acquired data—intrinsic activity, in the terminology of neuroscience. In fact, in order not to waste that computational capacity, computations related to learning and planning *should* be launched. That will enable the agent to act more intelligently when the time to act comes. This is also presumably why evolution has programmed wandering thoughts in us.<sup>11</sup>

Next, I consider in detail two different ways in which wandering thoughts can help in computation. In the first, the system is *planning future* actions by internally simulating the world, and trying out different, new actions to see which works best.<sup>12</sup> The second one is called *experience replay* because the system internally repeats memories of past behaviours and events exactly as they were perceived, in order to enable an iterative algorithm to efficiently learn from them.<sup>13</sup> In fact, a lot of what people simply call “thinking” falls into these two categories: You plan what to do in the future, and recall what happened to you in the past.<sup>14</sup>

### *Planning the future*

It is perhaps obvious why thinking about future actions is useful, as far as it is a case of planning. You can go through different kinds of plans and *simulate*, using your model of the world, what the results of your actions will be, and finally, choose the best one. In the case where you think about your job interview to take place tomorrow, you polish your answers beforehand by simulating what kind of impressions different options will make, eventually memorizing the best ones. Often, such thinking and planning is actually completely voluntary. If you really want to spend some time and energy to elaborate the best course of action, this is quite normal planning activity. When we talk about wandering thoughts, we mean a case where you consciously try to do something other than planning, but unrelated thinking nevertheless appears. It is the unwanted, intrusive quality of wandering thoughts that distinguishes them from ordinary thinking.<sup>15</sup>

You might actually want to relax and read a novel, but thoughts simulating the job interview just pop into your mind. This is understandable since as I just argued, it is especially during moments where you or the AI have nothing pressing to do that it would be a good idea (from the viewpoint of the designer of the system) to use the computing capacity for such planning. As we saw in Chapter 3, planning paths grow exponentially as a function of time, so there is a real need for using a lot of computation for planning.

The planning during wandering thoughts is a bit special in that it sometimes has no particular goal. It may be just looking at possible future paths in a big search tree to see what could be done to obtain rewards: a kind of ongoing, free-style planning. Such a search could actually be done by the Monte Carlo Tree Search algorithms

<sup>11</sup>Humans also sleep and dream: It is possible that the function of dreams is pretty much the same as that of wandering thoughts (Fox et al., 2013).

<sup>12</sup>(Baird et al., 2011)

<sup>13</sup>(Lin, 1991)

<sup>14</sup>The computations in planning are actually not that different from the computations in reinforcement learning, as was already discussed in Chapter 5. See also footnote 19 below on how the two computations could be combined, as well as a framework proposing something between these two kinds of action selection by Lengyel and Dayan (2008).

<sup>15</sup>In humans, planning can thus be done in various modes which differ in their relation to conscious control: there is controlled, consciously initiated planning on the one hand, and spontaneous, unconsciously initiated planning on the other hand. The spontaneous planning can further be divided to unwanted/wandering thoughts and simply spontaneous thoughts which are not unwanted (perhaps because you are lying on the sofa), see footnote 2.

(discussed in Chapter 7): they are randomly searching for plans, while focusing more on branches which seem to be more rewarding. It's a bit like thinking about what to do during the weekend when you're supposed to be concentrating on your work. Sometimes wandering thoughts do focus on planning for a specific goal: One theory proposes that wandering thoughts focus on goals that have been selected but not yet reached.<sup>16</sup> Typically, when you have been thinking about a difficult problem for a while without finding a solution, it will be difficult to relax and think about anything else, since that problem will constantly intrude on your mind.<sup>17</sup>

### *Experience replay for learning value functions*

In contrast to planning, it may be more difficult to understand why any system would like to simply repeat past experiences. You already saw what happened yesterday, so why repeat it in your mind, and why so many times? The reason is in the structure of the algorithms used in learning.

As we saw in Chapter 4, modern AI systems are based on learning from the data by using iterative algorithms. We saw the general idea of stochastic gradient methods: the data points (e.g. images) are presented to the system one by one, and a huge number of repetitions is needed. Most reinforcement algorithms are not, strictly speaking, stochastic gradient descent methods, but are closely related and share those properties. They proceed by observing the state of the world both before and after each action, as well as any reward obtained or punishment received. There are thus four pieces of information in what we might call a single “data point”: the state before the action, the state after the action, the action taken, and the reinforcement. Based on these, the system updates the state-value function.

What is crucial here is that, again, learning proceeds by making tiny modifications to the parameters of the system, in this case those computing the state-value function. Successful learning usually requires a huge number of iterations, or presentations of such actions and their consequences to the learning system. If you have access to really large amounts of data, you may just present each data point once, and learning will be

---

<sup>16</sup>(Klinger, 2013). His theory actually considers “spontaneous thoughts” which are more general than wandering thoughts. It further includes the interesting idea that wandering thoughts may not be just triggered when the computational capacity would be idle otherwise, but they could also be triggered when there is a goal, perhaps with an intention or commitment to it, but it is not currently possible to actually perform any meaningful action to reach the goal. Then, planning to reach that goal may be triggered involuntarily and lead to wandering thoughts.

<sup>17</sup>In game-playing AI, planning by simulation has been used in an extreme way in terms of “self-play”. A much-publicized example is AlphaGo, the system that first beat humans in the board game of Go, which we used as an example of dual processes earlier (Silver et al., 2016). After being input information on a huge number of actual games played by humans, it started playing against itself. This is a very special kind of planning, where you are simulating your opponent as part of the environment. Actually, there is no distinction between the agent itself and the opponent since the same agent “plays” both of them, and learns from the successes and failures of both of them. A later version of the AlphaGo system actually omits the learning from human games altogether and learns entirely by playing against itself; the ensuing system is aptly named AlphaGo Zero (Silver et al., 2017). Pure self-play has also allowed for an AI to rapidly approach human level in a highly complicated, multi-player esports video game called Dota 2: Using more than 100,000 processors running self-play in parallel, the OpenAI Five system can simulate in one day the same amount of data that would take more than a hundred years to collect in ordinary play against humans (OpenAI, 2018). Self-play was also used to achieve super-human performance in the game of poker (Brown and Sandholm, 2018). Actually, such learning by self-play was successfully used earlier in simpler games such as backgammon (Tesauro, 1995) and, even back in 1959, in checkers (Samuel, 1959), in one of the earliest machine learning projects. While some human knowledge was input to the learning process in most of the preceding studies, Tesauro (1995) also reported a variant with pure self-play similar to AlphaGo Zero. Something akin to self-play is actually used by humans when they are simulating social encounters in their own minds: We might use the same model for the actions of other people and the actions of ourselves, and learn both simultaneously. (Yet, the connection between our model of our own mind and our model of the minds of others is complex, see Carruthers (2009) for different possibilities.)

successful since the algorithm will have enough iterations anyway. However, the amount of data is typically limited. In the case of reinforcement learning, what is particularly slow is that the agent may need to act in a real environment and observe the consequences of its actions to gather data. One action by a robot can take a second or so, which is extremely slow compared to the processing speed of most computers and the potential speed of learning. Likewise, humans do not collect new experiences on, say, job interviews, that often.<sup>18</sup>

This is where experience replay is useful. It means that in reinforcement learning, the system is not just using the data related to the most recent action and then throwing it away; instead, it stores the data, and re-uses past actions and the states associated with them many times. That is, it “replays” or recalls past actions and events and uses them in the iterative learning algorithm as if they happened now. This improves the performance of the learning algorithm by enabling it to make many more iterations with each data point, and thus many more iterations for the same limited amount of data. This is how more information is extracted from the data. Another great utility is that usually the agent can retrieve past events from memory much faster than actually acting in the world, and thus replay enables learning much faster. There are other reasons as well, as we will see below.

Obviously, there is a trade-off here: If you just use all your time replaying old events from memory, you will not get new data about reinforcement resulting from actions. So, you cannot use all your time for just replay. It should be smart to engage in replay when the environment does not enable too many meaningful actions — in plain English, when nothing interesting is happening and the agent is “bored”, which points directly at wandering thoughts.

It is also possible to do something between pure replay and planning. You can replay past events while trying out different actions in a simulation. This means the system starts by recalling something that happened earlier, but then it simulates what would have happened if it had acted differently. Certainly, we all have experienced such wandering thoughts: “If, yesterday, in that situation, I had done X instead of Y...” This is even better than just replaying actual past events since the system is then creating *new* data using past events together with its model of the world.<sup>19</sup>

## Experience replay focuses on reinforcing events

Any replay method must choose *which* events, or short “episodes” of events, it will replay. A system that has gathered a lot of data on past actions cannot replay all its history. Likewise for planning: if the system starts planning in its idle time, it needs to choose the starting state for its plan—what kind of situation does your fantasy start in—and perhaps a goal as well.

A dominant idea in AI is that replay should prioritize events where any kind of reinforcement signal was obtained, whether positive or negative, and this seems to be the case in the brain as well. Experienced episodes containing such events are the most important in computing the state-value function. This may help explain why we have so many wandering thoughts about negative events. When you do something embarrassing, it

<sup>18</sup>In fact, calculations of the amount of data that humans observe in real life show that the number of data points needed in AI is often much larger than what humans need. For example, children seem to learn to speak from a relatively small number of “input” words (Cristia et al., 2019); current AI systems need orders of magnitude more. Perhaps even more strikingly, humans can learn from a single example they see or hear, as I have pointed out earlier, see e.g. Lake et al. (2015). This shows that there is a lot of room for improvement in AI compared to the brain in terms of efficiently using all the data.

<sup>19</sup>(Sutton, 1991; Sutton and Barto, 2018). See also Mattar and Daw (2018) for a theoretical unification of replay and planning.



may replay in your mind many, many times. This should be useful so that you learn to associate the negative reinforcement (social embarrassment) with the actions you took in that particular situation, thus improving your estimate of the state-value function—and future behaviour. A similar mechanism might be at play when choosing the starting states for planning, although current theory does not explicitly explain that.

It has been found that replay of past events can be particularly useful if the experience is replayed backwards, starting from reinforcing events. Suppose a robot gets a particularly nice reward (say, a lot of energy in its batteries) whenever it finds itself in room #42 of a building where it cleans the floors. Based on this experience alone, it will immediately assign a large state-value to room #42. But in order to find room #42 in the future, it has to code its location with respect to the other rooms in the state-value function. This is easy to do if it replays its path to room #42 in reverse order. Suppose just before arriving in room #42 it was in room #13, and before that, room #21. It replays the sequence in reverse: #42, #13, #21. Now, it will assign large but slightly decreasing state-values to each of these rooms, so that the state-value is decreasing the further the replay goes—the decreases are justified by the theory of discounting. The end result is that while #42 has the largest state-value, #13 has a rather large one as well, and #21 is not far behind. Now, if the agent ever finds itself again next to room #21, it knows that to find a state with a large state-value, it should enter room #21, and there it will understand the best choice for the next state is #13, and eventually #42. (It may sound like all this could be learned by a single replay, but in reality it must happen by smaller increments to properly combine information from many different paths and data points.) Combining such backward replay with the above-mentioned prioritization of reinforcing events leads to a method called “prioritized sweeping”.<sup>20</sup>

If wandering thoughts use such a prioritizing form of replay, we see that they are closely related to the theory of emotions as interrupts discussed in Chapter 8. Both mechanisms direct the agent’s processing (one might say attention) towards dangerous or rewarding events. Emotional interrupts are more primitive, typically focused on easily identifiable and evolutionarily important threats which are present in the current state. In contrast, wandering thoughts are about learning when no threat is currently observed, potentially leading to quite sophisticated behaviours.<sup>21</sup>

---

<sup>20</sup> (Moore and Atkeson, 1993; Schaul et al., 2016; Singer and Frank, 2009) More precisely, the prioritization mechanism replays memories of individual states (and actions taken in them) whose replay leads to maximal change in the estimated state-value function. This is not exactly the same as replaying episodes where a strong reinforcement occurred, as proposed earlier in the text, but it is closely related. Typically, a strong reward or punishment is unexpected, at least in the beginning of the learning. When you find a reward the first time, your state-value function is in some rather random initial state, and you could not really predict that the reward would be obtained; thus any reward is initially surprising. That is why prioritized sweeping prioritizes, as a first approximation, episodes containing reward or punishment.

<sup>21</sup> A lot of replay is probably related to rewards, and thus to planning and RL, but some part of wandering thoughts and replay is clearly independent of any rewards. We saw earlier that people are able to perform unsupervised or supervised learning from a single representation of a data point (page 75). If you hear a nice melody, it may be replayed in your mind repetitively, even quite obsessively. Such replay is best understood as performing some kind of unsupervised learning—which does not need any kind of reward or reinforcement signal. For example, it can be Hebbian learning or some kind of feature extraction, which learns the melody and its characteristics particularly well by repetition. The crucial similarity between reinforcement learning, Hebbian learning, and most kinds of machine learning is their iterative nature, and in particular, the need for many iterations. Some of that data may not be real data replayed, but simulated data more akin to planning; such simulation can in fact be used to perform learning in a Bayesian framework (Gutmann et al., 2018). An alternative theory on resting-state activity actually links it to the priors used in Bayesian perception (Berkes et al., 2011; Aitchison and Lengyel, 2016; Hoyer and Hyvärinen, 2003). The idea is that activities of the neurons in resting-state, at least in the sensory cortices, follow the prior distribution of those features that they are encoding. While this theory is not framed in terms of replay, we could interpret it as saying that resting-state activity is in some sense “replaying” typical sensory inputs. These two theories may thus not be incompatible, the replay or wandering thoughts theory focusing on reward processing and the Bayesian

## Replay exists in rats, humans, and machines

Replay has long been observed in neuroscience experiments. Typical experiments measure brain activity in rats which are running in a maze, seeking food or drink. A brain area called the hippocampus is specialized in storing episodes and events —such as the sequence of running forward, turning left or right, and finding cheese. It is thought that the hippocampus replays such episodes, simultaneously signalling them to other brain areas, which then use such replayed input for learning. Replay was initially observed during sleep, but it can also be seen in awake rats.<sup>22</sup> Recent experiments also show that something similar to prioritized sweeping, where the events are replayed backwards, seems to be happening in rodents.<sup>23</sup>

Research has also found brain activations that look like planning: a rat can initiate sequences of events which it has not yet experienced, but which it might perform in the future. For example, the rat can in some sense “imagine” a possible trajectory in a maze, which it may or may not follow later.<sup>24</sup> So, the mammalian brain seems to use strategies which are very similar to what you would expect from the design considerations of AI. This is not surprising since the brain and AI are trying to solve the same computational problems; but it is also the case because the AI designs have been influenced by our knowledge of what happens in the brain.

It may in fact be that such processing in rats is not very different from wandering thoughts considered in human psychology. Something at least resembling replay by prioritized sweeping can also be observed in the human brain, although the limitations in measurement technology make it difficult to draw exact parallels.<sup>25</sup> While replay is usually connected with the hippocampus, and planning with the default-mode network, the hippocampus is actually part of the default-mode network according to some definitions.<sup>26</sup> (Rats do have a default-mode network just like humans.<sup>27</sup>) The connection between wandering thoughts and the hippocampus is also seen in the fact that people with damage in the hippocampus have difficulties in imagining new experiences.<sup>28</sup>

Some scientists are reluctant to make such parallels between hippocampal replay and wandering thoughts, since they would seem to imply that rats “think” or “imagine” like humans, at least in the sense that rats engage in planning by imagining different sequences of actions and choose the best one.<sup>29</sup> Likewise, we immediately run into the question of whether such replay in an AI means that we would have to admit that an AI can “think”. “Thinking” is not a well-defined concept in either neuroscience or AI, which makes this question difficult to answer.<sup>30</sup>

---

theory focusing on basic sensory processing.

<sup>22</sup>Sleep: (Buzsáki, 1996), awake animals: (Karlsson and Frank, 2009)

<sup>23</sup>(Diba and Buzsáki, 2007; Ambrose et al., 2016)

<sup>24</sup>This is called *preplay* in neuroscience (Pfeiffer and Foster, 2013; Wikenheiser and Redish, 2015). Some forms of replay or preplay seem to be happening much faster than real-time, which would make it particularly useful (Karlsson and Frank, 2009).

<sup>25</sup>(Buckner, 2010; Gruber et al., 2016; Kurth-Nelson et al., 2016; Momennejad et al., 2018)

<sup>26</sup>(Andrews-Hanna, 2012)

<sup>27</sup>(Lu et al., 2012)

<sup>28</sup>(Hassabis et al., 2007)

<sup>29</sup>See (Redshaw and Bulley, 2018; Corballis, 2019) for discussions on whether non-human animals might possess such capabilities.

<sup>30</sup>An important point is that, sometimes, thinking is defined to be conscious, and humans do a lot of the simulation consciously, while the AI is probably not conscious when planning. However, I see no reason to make a connection between consciousness on the one hand, and the phenomena of replay, planning, or imagining future actions on the other. Consciousness is a separate topic to which we will return in Chapter 12.

*Creative thinking*

The discussion so far considers wandering thoughts as rather mechanistic solutions to some well-defined computational problems. This does not do justice to the variety of wandering thoughts in humans. Spontaneous thinking can be tremendously creative; in fact, it is one of the critical aspects of human creativity.

Now, what is creativity? As a first approach, we might actually think of planning as a creative activity. You have the current state, a goal, and you have to somehow create a path between the two. In fact, many different kinds of problem-solving could be seen as special cases of such planning: even proving a mathematical theorem can be formalized as planning a “route” from the premises to the conclusion of the theorem. However, some would argue that this is just running an algorithm, so it cannot be called creative. I wonder why running an algorithm could not be called creative. What else does an intelligent system do anyway? On a sufficiently high level of abstraction, is not all our thinking a product of various kinds of algorithms? I shall not attempt to answer the question of what creativity is here; I will just note that creativity is not easy to define; it is in that sense very similar to the concept of intelligence.

I think a randomized algorithm, such as Monte Carlo Tree Search, could be quite convincing as an example of creativity. Such algorithms contain certain randomness in their computation which makes the algorithm try out completely new paths or ideas. They are not just deterministically finding a single solution to a given problem, but rather creatively imagining, as it were, a number of possible things to do, or steps towards a solution to the problem. As we saw above, such randomized algorithms have been very successful in game-playing AI, and they also offer a plausible model for some of the wandering thoughts. From this viewpoint, it is natural that the computations performed by wandering thoughts can also result in creative problem-solving.<sup>31</sup>

In fact, there are also some wandering thoughts that cannot be plausibly considered as replay or planning. Perhaps, while lying idly on your sofa, you have a series of seemingly unrelated mental images, or a superhero fantasy that could never actually happen in reality. One function of such wandering thoughts may be to create completely new ideas and associations, even new goals. In AI approaches to creativity, a “generate and test” approach is wide-spread: it means that new items are more or less randomly generated by one part of the system, and then another part of the system tests whether they make any sense. Unrealistic, weird, and unstructured wandering thoughts could be the result of such random generation; hopefully, our more rational part then tests them and decides which ones make any sense and should be taken seriously.<sup>32</sup>

## Wandering thoughts multiply suffering

So far, we have seen that while mind wandering may be detrimental for whatever you’re trying to do at the present moment, it helps in planning and learning, perhaps even allowing some creativity. From a purely information-processing viewpoint, it is probably a useful thing since similar ideas are currently used in AI systems, and after all, evolution would not have “programmed” us to have a wandering mind if it were not useful to us from the evolutionary viewpoint.

Yet, evolution does not try to make us happy. A problem with replaying past memories and planning the future in human brains is that we are, on some level, unable to understand they are not real. If you remember

---

<sup>31</sup> (Fox and Beaty, 2019)

<sup>32</sup> For surveys on what is called “computational creativity”, i.e. trying to make computers creative, see Colton et al. (2009); Toivonen and Gross (2015).

an embarrassing episode from the past, you actually feel embarrassed. If you think about something scary that might happen to you tomorrow, you actually start feeling scared. That is, wandering thoughts increase human suffering by making us suffer from simulated or replayed events, in addition to the real ones.

Any suffering produced by real-life events may, in fact, be repeated *many times* by the replay of those events. Likewise, if something unpleasant is expected to happen, the unpleasantness, the threat, is felt many times in planning how to avoid that thing—which may actually turn out not to happen at all. Planning future events even includes frustration when things in the fantasy don't go as you would like them to, and you can be frustrated many times by the planning of a single event. Due to this multiplication of suffering by wandering thoughts, it could be argued that the *vast majority* of our suffering actually comes from remembering or anticipating unpleasant events. The anticipation is closely related to what we said about fear in the preceding chapters, but the aspect of replaying unpleasant memories is new.

Importantly from the viewpoint of suffering, you have little control regarding such wandering thoughts. You may think that *you* must have decided to recall an unpleasant conversation, but in fact, the recollection and replay just started without you deciding anything, and even if you try to think about something else, you may find yourself unable to do so. This is another clear connection to the emotional interrupts: both wandering thoughts and emotional interrupts are largely beyond conscious control. You cannot switch off those systems. In fact, it is even worse: both systems actually take control of the whole agent.

Some research actually claims that the wandering mind is generally unhappy. That sounds plausible if wandering thoughts multiply suffering, as I just argued. However, it might be a bit of an overgeneralization.<sup>33</sup> Whether wandering thoughts make you unhappy probably depends on their contents. It might seem obvious that having wandering thoughts with negative feelings, such as worrying, makes you unhappier, while positive content has the opposite effect. In what is a rather extreme case, a study found that women having wandering thoughts about their significant others actually felt happier.<sup>34</sup> Close to the negative extreme, we find *rumination*, which is thinking about negative events that typically happened in the past and are related to one's personal concerns.<sup>35</sup> It is particularly frequent in depression and, unsurprisingly, leads to low mood. For individuals with depressive tendencies, most wandering thoughts may consist of depressive rumination, and eventually may lead to relapse and full-blown depressive episodes; even for normal individuals, wandering thoughts provide an opportunity for rumination to arise, and thus may lead, on the average, to negative mood.<sup>36</sup> In spite of some reservations, therefore, I think an important point is made in claiming that a wan-

<sup>33</sup>See Killingsworth and Gilbert (2010) for the claim that "A wandering mind is an unhappy mind". The problem is, however, that such studies don't usually show that it is mind-wandering that *makes* people unhappy. It is also possible that the causal effect is the opposite: when we are unhappy, thoughts start wandering more (Smallwood et al., 2009). This could be because negative mood is related to unresolved goals or personal problems, which are then processed during mind-wandering. If you're sad, it may be because you are experiencing problems, and those problems need extra processing by mind wandering (Poerio et al., 2013), in line with footnote 16 above.

<sup>34</sup>(Poerio et al., 2015). Intriguingly, the effect of wandering thoughts on mood may strongly depend on whether you think about the past or the future. Ruby et al. (2013) found that future-oriented thinking has a general positive effect on mood, even if the contents were negative; perhaps this is so because when we solve a planning problem, we get happier. In contrast, thinking about past events was found to lead to negative mood independently of the contents of the thoughts.

<sup>35</sup>(Whitmer and Gotlib, 2013). Perhaps the most extreme negative example would be flashbacks about a traumatizing event in post-traumatic stress disorder (Yehuda, 2002).

<sup>36</sup>(Teasdale et al., 2000; Marchetti et al., 2014, 2016; Ottaviani et al., 2013; Van Vugt et al., 2018). Another point to note here is that when talking about thoughts inducing a positive or negative mood, we need a baseline. In the research cited, this is typically a rather normal, average mood. However, if the baseline had absolutely no wandering thoughts, as might be considered ideal in some particular

dering mind is an unhappy mind; we will get back to this important point when talking about meditation in Chapter 15.

*Why do wandering thoughts trigger feelings?*

Replaying negative experiences, or planning the future, might not have anything to do with suffering if they did not somehow *feel* unpleasant, i.e. if they did not activate the negative valence signalling. A person may have reoccurring wandering thoughts about going to the dentist and vaguely feel the pain that the dentist's tools will cause in her mouth. Isn't it odd that she feels the pain although she is not at the dentist at all? While you probably have to go to the dentist one day, people also worry about the possibility of various disasters that are not at all likely to happen to them. Let me repeat Montaigne's comment: "One who fears suffering is already suffering from what he fears".

Thoughts rarely correspond to something that is actually happening here and now, as opposed to perceptions. Almost by definition, our thinking is usually about past events which are no longer there, or future events which have not yet happened, and may not happen at all. Why do we then feel upset about them, or, from a more computational viewpoint, why do they activate negative valence signals? Indeed—this is a deep question that we encounter several times in this book—why do we feel the emotions associated with memories and imagination?

From the viewpoint of computational design, it is clear that the system that computes state-values and predicts rewards has to be active in wandering thoughts, at least to some extent, so that the brain can take its evaluations into account when planning and learning. What does not seem necessary is that we actually, on a visceral level, feel pleasant or unpleasant about the events produced by planned actions. Why do our bodies react to our fantasies as if they were true? I suggest this is a kind of a *computational shortcut*. If you want to make learning from the simulation as simple as possible, it makes sense to use the same mechanisms and networks as in the case of real data. This is possible if the AI or the brain doing the reinforcement learning is fed the same kind of signals and into the same networks regardless of whether the action is real or simulated.

Ultimately, combined with the hypothesis that the error signals are best broadcast to the whole brain using the pain system (Chapter 2), such computational simplification seems to have led to a situation where in the brain, it is not possible to give an "unpleasant" signal to the planning system without activating the main system that signals suffering to the whole system. In other words, perhaps humans *feel* suffering during negative wandering thoughts simply because it makes the design of the learning system easier.<sup>37</sup>

Here we see a particularly striking conflict between evolutionary goals and happiness. Suffering from the simulation of negative events may be a computational shortcut, which is not really that necessary. It is just that the brain was "designed" by evolutionary forces which do not care if the system design makes you suffer many times more; they just found this design to be functional for their own evolutionary purposes.

---

meditation traditions, it could be that even positively valenced wandering thoughts actually have a negative effect on the mood.

<sup>37</sup>Incidentally, such processing is also part of the somatic marker hypothesis; see footnote 24 in Chapter 8.

## Chapter 10

# Perception as construction of the world

Without any perceptual abilities, an agent can hardly do anything intelligent in the real world. Neural networks already give a rudimentary system for perception: for an input image, they can try to tell what it depicts. But it turns out that perception is an extremely difficult problem. In this chapter, I explain the main difficulties involved in perception, and how they can be, to some very limited extent, solved by modern AI. However, I argue that the very problem of perception is extremely difficult, and even our brains do not solve it very well. Here, I consider in detail visual perception, but the theory largely holds for other kinds of perception, such as auditory perception.

What is crucial for the main theme of this book is to understand the implications of the extreme difficulty of perception. It means our perceptions are quite *uncertain*, or unreliable, and much more so than we tend to realize. We fill in the gaps in the incoming information by using various assumptions, or prior information, about the world. Therefore, one aspect of such uncertainty is subjectivity: we fill in the gaps using our *own* assumptions, and my assumptions may be different from yours. Perception is essentially a construction, a result of unreliable and somewhat arbitrary computations; it is not an objective and perfect recovery of some underlying truth.

These fundamental problems in perception feed into the difficulty of making correct inferences about the world: they make any categorization uncertain, they reduce the possibility of predicting the world, and consequently reduce any control the agent has. This increases various errors such as reward prediction errors, and thus suffering. More specifically, computation of reward loss is dependent on the prediction of reward as well as the perception of obtained reward, which are both subject to the limitations of perception, and thus it can go wrong.

### Vision only seems to be effortless and certain

It may be suprising to many people to see how difficult computer vision actually is, and what an incredible feat the visual system of our brain is accomplishing, literally, every second. It all seems to happen so effortlessly and automatically. However, our capacity for vision is effortless only in the sense that it does not require much *conscious* effort, and it is automatic only in the sense that it does not usually need any *conscious* decisions or thinking. You turn your gaze towards a cat, and immediately, without any conscious effort, you recognize it as a cat. This is a typical, even extreme case of dual systems processing: most of the computations happen in

neural networks, not at the level of symbolic, conscious thinking. Since we have little access to the processing in the neural networks, we cannot understand how complicated their computations are.

In the early days of AI in the 1970s, computer scientists thought programming such “computer vision” must be easy. However, anybody either studying the human visual system or trying to build a computer vision system is quickly convinced of the near-miraculous complexity of the information-processing needed for vision, and performed by our brain almost all the time. Knowing that history, it is not surprising that while computers can beat humans in chess or arithmetics, they are nowhere near human performance in visual processing.<sup>1</sup>

### *Too much data*

A major difficulty in vision is the huge amount of data received by the system. The immensity of the data is perhaps obvious to anybody who has waited for video data to download over a mediocre internet connection. In fact, the vast majority of internet traffic takes the form of video data. Text data is completely negligible in terms of file size: a large book is hardly equal to a second of video data.

Likewise, humans and other mammals receive a huge, continuous stream of data from the environment through their eyes. The human retina contains something like one hundred million photoreceptors, which are cells that convert incoming light into neural signals. The manner in which the data is stored and transmitted may be very different from computers, but still the fundamental problem of receiving an immense amount of data is there, as well as the requirement of a huge amount of information-processing capacity. In fact, the visual areas constitute something like half of the human cerebral cortex—the part of the brain where most sophisticated processing takes place.<sup>2</sup>

### *Yet, information is missing*

Having such huge amounts of data is both a blessing and a curse. A curse obviously in the sense that you need immense computing power to handle such a data deluge; a blessing in the sense that such huge amounts of data may contain a lot of useful information. Yet, paradoxically, the information contained in the input to a camera or the retina is almost always lacking in various important ways.

One of the most fundamental problems in vision is that what each eye gives us is a two-dimensional projection of the world, just like an ordinary photograph. A photograph is nothing like a 3D hologram: most of the information on the 3D structure of the objects is missing. (Having two eyes gives some hints of the 3D structure, i.e. which objects are close to you and which are far-away, but this only slightly remedies the problem.)

Suppose you see a black cat. Now, the actual 2D projection will be very different when you can see the cat from different viewpoints: from the front, from one side, from the other side, from above, and so on. That is, the pixels which are black are not at all the same in the different cases; the pixel values that would be input to a neural network will vary widely when the cat is seen from different viewpoints. Thus, the neural network will have to somehow understand that very different pixel values correspond to the same object. To illustrate this problem of 3D to 2D conversion, consider what even a simple cube can look like in different projections.

---

<sup>1</sup> It is true that in some specific, well-defined tasks, such as recognizing animals in photographs, AI can actually outperform humans. However, such performance is usually specific to a certain kind of input data and task, and it is still far away from the versatility of human vision; see e.g. discussion by Recht et al. (2019).

<sup>2</sup> (Nakayama, 1999)

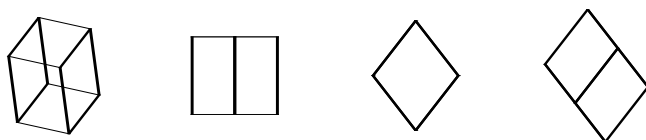


Figure 10.1: An illustration of the inverse problem that makes vision particularly difficult. The four figures are all 2D projections of the same 3D cube. Any camera or a single eye can only capture one such 2D projection, which loses a lot of information and creates ambiguity.

Some possibilities are shown in Figure 10.1. Its 2D projection can look like a rectangle (possibly a square), like a diamond, and many other things.

And this is just one kind of information lacking. More fundamentally, the problem is that any object can undergo many different kinds of *transformations*. Consider a cat again: it can take many different shapes by moving its limbs; sometimes its legs are wide apart, sometimes close to each other. Sometimes it stretches its whole body, sometimes it puffs up. If you think about the 2D image created, it will again be quite different in these different cases. As another example, the lighting conditions can be very different. Imagine that light comes from above, or from behind the cat: Again the cat looks very different, and even more so in a 2D projection.

Those were some of the problems in recognizing a single cat. To make things even more complicated, different cats look very different. Some are black, some are white, so the pixel values are even more fundamentally different. Yet, you somehow are able to see that they are all cats.

Such ambiguity, or incompleteness of visual information in a camera or the retina is the reason why vision is called an *inverse problem*.<sup>3</sup> As a very simple illustration of an inverse problem, consider there are two numbers which we denote by the variables  $x$  and  $y$ . You want to know both these numbers, but the trick is you only are given their sum,  $x + y$ . How could you possibly find out both of those original numbers—how can you “invert” the equation? Suppose you are told the sum of two numbers is equal to 10. There are many possibilities what the actual  $x$  and  $y$  may be like, for example,  $x = 5$  and  $y = 5$ , or  $x = 7$  and  $y = 3$  etc. Vision is a lot like this. What you observe are the pixel values in, say, a photograph. But there are a lot of factors that determine what the pixel values are like: the identity of the object in the photograph, the location of the object, the lighting conditions, the background, to name just a few. It is next to impossible to figure out what there is in the image without some tricks.

Actually, the fact that perception is based on incomplete information is in some sense quite blatant. Just think about the fact that you cannot see through solid surfaces. Suppose you look at a wall in front of you: you cannot see what is on the other side. Your perception is limited by the physics of light, which does not penetrate the wall, and thus you only obtain limited data and limited information about the environment. That may be an extreme example, but the point is that all perception is similarly constrained; it is just a matter of degree. Curiously, in your mind, you do have some idea of what there is behind the wall (another room, the street, or something else), but this idea is vague and uncertain. We will see in this chapter why all perception is, to some extent, a similar kind of guesswork.

<sup>3</sup>Strictly speaking, what we consider here is an *ill-posed* inverse problem; however, ill-posedness is often implicitly assumed when talking about inverse problems.



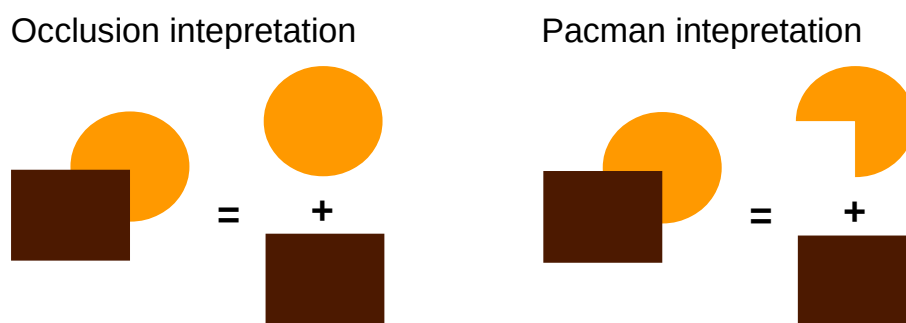


Figure 10.2: Suppose you see the figure on the far left, consisting of a square and a part of a disk. On the left: a typical interpretation where it is assumed that the disk is complete but occluded. On the right: another logically possible interpretation, with a pacman “eating” the square, but one that our visual system would not make because it is less likely. Our visual system chooses the interpretation which is more likely, given its prior information about the environment.

## Perception as unconscious inference

Yet, AI has recently been making major progress in vision. One reason is that computers have been getting much faster every year, but that is of course not enough in itself if you don’t know how to program your computer. The crucial breakthrough in recent computer vision has been the application of neural networks. Neural networks offer two important advances. First, they enable the processing of vast amounts of data to be distributed into a large number of processors, which work in parallel and thus can process the data more easily. The advantages of such distributed and parallel processing are considered in more detail in Chapter 11. In this chapter, we focus on another advantage, which is that we know how to use neural networks to *learn* from big data sets. Learning can alleviate, and to some extent solve, the problem of incomplete information, such as seeing only a 2D projection of the 3D world.

The trick is to learn what the world *typically* looks like, and to use the learned regularities to complement the incoming data. Look at the figure on the far left-hand side of Figure 10.2. Here, we tend to perceive a disk and a square. This is because we immediately assume that the disk actually continues behind the square, it is just partly occluded (i.e. blocked from view) by the rectangle. In fact, we tend to almost *see* a whole disk. There’s nothing wrong with such an assumption, but it does not necessarily follow from the figure. Alternative interpretations are possible based on this incomplete data. For example, it could be that the figure actually consists of a square and a “pacman”, as illustrated on the right-hand side of the figure.

Perceptions such as in this example are usually explained as results of *unconscious inference using prior information*. The visual system has learned certain regularities in the outside world—this is called prior information. For example, contours are typically continuous and smooth; lines are typically long and straight; objects can be behind or in front of each other. So, in Figure 10.2, the brain computes that it is very likely that the incomplete disk is actually part of a whole disk, but we just don’t get visual input on the whole disk because it is blocked. Such a conclusion is made by neural networks which are outside of our consciousness, thus the process is called unconscious inference.<sup>4</sup>

<sup>4</sup>Inference means the computational process leading to a conclusion or a decision.

The inference in question is also *probabilistic*: The visual system cannot know for sure whether the edge of the disk continues behind the square, but it is more likely that it does than that it doesn't. That is, the brain cannot make any judgements that are logically necessary and certain about this picture. The only thing the brain can do is to calculate probabilities and choose the most probable interpretation.<sup>5</sup> That is why perception is necessarily *uncertain*.

### *Bayesian inference*

The probabilistic inference needed in perception takes a particular form where the goal is to determine causes when observing the effects. The mathematical theory behind such inference was initially proposed by Thomas Bayes in the 18th century, which is why such inference is often called Bayesian.<sup>6</sup> In the case of perception, the “effects” are the patterns of light coming into your eyes, while the “causes” are the objects and events in the outside world.

Typical scientific models based on physics will tell you what the effects are for given causes. For example, given an object and its location in your field of vision, you can rather easily compute, by basic physics, what the light coming from that object to your eyes will be like. But doing the computation backwards is more difficult. Given that your eye receives certain light patterns, as registered by your sensory organs, how can you know what went on in the outside world? You have to somehow *invert* your physical model of the world, leading to the inverse problems just mentioned. Such problems can be approached by Bayesian inference, especially in the case where we can only calculate probabilities, which is exactly the case here.

Bayesian inference tells that the probability for a given cause (given we observe certain effects) is proportional to the product of two things: First, the probability that such a cause creates the observed effects, and second, how likely the cause is to occur in general. The first part here is rather obvious: A given cause is more likely to be responsible for what your sensory organs report if the cause and such sensory input are compatible, and that cause is likely to produce the observed effects. However, the important point here is in the second part: A given cause is even more probable if its general probability of occurrence is large. That is, if the cause has high “prior probability” in the terminology of Bayesian inference.<sup>7</sup>

<sup>5</sup>A fundamental question is whether the brain chooses one interpretation or whether it can entertain several interpretations simultaneously. Something in between these two seems to be happening in the special case of bistable perception, which means that when a stimulus can very well be interpreted in two different ways, the two interpretations seems to be alternating in the brain, so that conscious perception switches from one interpretation to another every few seconds or so (Sterzer et al., 2009). The proportion of time allocated to each interpretation may, in fact, reflect its probability that the brain computes by Bayesian inference discussed next in the text (Moreno-Bote et al., 2011).

<sup>6</sup>For neuroscience-oriented introductions, see (Kersten et al., 2004; Ma et al., 2022). While Thomas Bayes is usually credited with the general mathematical theory used in this context, the specific idea of perception as unconscious inference was actually formulated later by Hermann von Helmholtz, which is why some authors call this framework the Helmholtzian theory of perception. (Also, the credit for the mathematical theory should perhaps largely go to Pierre-Simon Laplace.) The recent proposal of a “free-energy” brain theory (Friston, 2010) is essentially a reformulation of these ideas, with some additional hypotheses extending it to action selection.

<sup>7</sup>To get into more mathematical detail, Bayesian inference wants to compute the probability  $P(\text{cause given effect})$ , where  $P$  denotes probability. More precisely, this is a conditional probability, i.e. the probability of one thing (cause) given that another thing (effect) has been observed. This is the typical case of inference: we observe the effects and want to find the causes, or at least their probabilities. The celebrated Bayes formula then says the aforementioned probability is equal to  $P(\text{effect given cause}) \times P(\text{cause}) / P(\text{effect})$ . Here, the term  $P(\text{effect given cause})$  can be computed from a physical model of the world implemented in your brain.  $P(\text{cause})$  is the prior probability of a given cause; this is where the prior information about what typically happens in the world comes in.  $P(\text{effect})$  is not so important because we are not comparing different effects, so it is constant, and it can actually be computed from the other probabilities

Consider the following example. Through your living room window, you get a glimpse of something green moving on the street. It could have been a green car, or it could have been a Martian (they are all green, as we all know). Both of these two causes (car or Martian) would produce the same kind of quick flash of something green moving on the street, or more precisely, some green light briefly entering your eyes. So, the probability of the effect (green light stimulating your retina) is high for both two causes; let's say for the sake of argument that it is equally high in both cases. However, you will not think it is a Martian. The reason is that your brain uses Bayesian inference and looks at the prior probabilities. The prior probability of a Martian is very much lower than the prior probability of a green car; the brain knows that in general, it is very rare to encounter any Martians. Thus, in weighing the probabilities of the different causes, the green car wins by a wide margin. This inference is possible because the brain has a model of what the world is typically, or probably, like: Martians are quite rarely encountered, at least on planet Earth.

## Prior information can be learned

Prior information, i.e. a model of what the world is typically like, is central in such unconscious inference, so where does it come from? The crucial principle in modern AI and neuroscience is that the prior information can be obtained by learning from data: Learning is thus the basis of perception. Now that may seem like a weird claim from a biological viewpoint. How could perception possibly be based on learning, given that many animals see quite well more or less immediately after birth? With human infants, developing proper vision actually takes several months but that is beside the point. The point here is to understand the different meanings of the word “learning”.

When I talk about learning here, I mean learning in a very abstract sense where a system *adapts* its behaviour and computations to the environment in which it operates, and in particular to the input it receives. In human perception, such adaptation happens on different levels and time scales: there is both the evolutionary adaptation and the development of the individual (after birth). These two time scales are very different, but if we are interested in the final result of learning, we can just lump the two kinds of adaptation together. Likewise, the optimization procedures are very different: evolution is based on natural selection while individual development presumably uses something like Hebbian learning—although we don't understand the details yet. Again, if we just look at the end result of the combination of those processes, we can ignore the difference of optimization procedures as well, and simply call this whole process “learning”. This resolves the paradox of animals being able to perceive things instantly after birth. Their sensory processing is using all the results of the evolutionary part of learning, and thus even before having received much input as individuals, their neural networks are capable of some rudimentary processing.<sup>8</sup>

---

by a simple formula.

<sup>8</sup>There is actually something in between those two kinds (evolutionary and developmental) of biological learning, which is learning in the womb. At the late stages of the pregnancy, the visual system of the foetus is “learning”. While its eyes are closed, and they don't receive much input, certain dynamic patterns called “travelling waves” are generated in the eye, on the retina. These patterns are then fed to the visual cortex of the brain, enabling some basic learning of visual regularities, complementing the information in the genes (Wong, 1999).

*Neural networks weights contain the prior information*

We already saw in Chapter 4 how it is possible to train a neural network from big data sets. The weights in the network are learned based on minimization of some error function. In the simplest case, the learning algorithm knows what there is in each image used for training (a cat or a dog) which provides a label or a category, and then we can use supervised learning. If we want to understand biological vision, though, unsupervised learning is preferred. This is because the visual system does not really have anybody constantly giving labels to each input image, which makes supervised learning unrealistic as a theoretical framework.

Fortunately, Hebbian learning and other methods of unsupervised learning can learn to analyse images in interesting ways, without any supervision. Intuitively, if the input to neuron A and the input to neuron B are often rather similar, it is likely that they are somehow signalling the same thing, and thus they should be processed together, for example by computing their average or difference.<sup>9</sup> The results of such learning are stored in the synaptic weights of the neural network. From the viewpoint of Bayesian perception, we can thus say that the prior information is learned and stored in the form of the weights connecting the neurons. This allows us to investigate what kind of prior information can be learned from visual input, by looking at the weights of an artificial neural network trained with ordinary photographs as input.

If we look at the initial analysis of images done by a neural network with one layer, different learning rules almost invariably give the same result: the most basic visual regularities are something like short edges or bars. Figure 10.3 shows some examples. Interestingly, such AI learning leads to processing which is very similar to the part of the brain that does some of the earliest analysis of incoming images, called the primary visual cortex. Measurements of many cells in that area reveal that they compute features which look very much like those in Fig 10.3. Such edges and bars are clearly a very fundamental property of the structure of images, which is why they are often used in the beginning of image analysis.<sup>10</sup>

Such edges and bars can be seen as the first stage of the successive “pattern-matching” on which neural network computation is based. We can actually go further and train a feedforward neural network with many layers to analyse images.<sup>11</sup> After successful training, a multi-layer neural network can contain extremely rich

<sup>9</sup>In particular, Hebbian learning can implement feature extraction methods such as principal component analysis (Oja, 1982) and independent component analysis (Hyvärinen et al., 2001).

<sup>10</sup>Based on (Olshausen and Field, 1996; Van Hateren and van der Schaaf, 1998). The learning principle used here can be intuitively understood from two different viewpoints. One is independence of the features: the outputs of the neural network (which in this case has a single layer) should be as independent as possible in the sense of probability theory. In other words, knowing one feature should give minimal information about the other features. The other viewpoint is sparsity: the features should be silent (zero) most of the time and only rarely turned “on”. An important benefit of such sparse coding is that it minimizes energy consumption if representing a feature that is zero consumes little energy. Therefore, the learning principle used is called either independent component analysis or sparse coding, which are almost the same thing. Such analysis can be implemented as a particular kind of Hebbian learning. Actually, there is an even more fundamental regularity in visual input than the one depicted here, which is that two near-by pixels tend to have similar grey-scale values (they are strongly correlated). That is, if a pixel is, say, white, the pixels next to it are quite likely to be white as well—and the same applies for any colour. Such similarities are analysed by neurons (“ganglion cells”) in the retina. However, this regularity is so elementary that it is in some sense included in, or implied by, the regularity described by the edges. Mathematically speaking, the covariances of pixel grey-scale values are perfectly modelled by independent component analysis and no additional model is needed. For a general introduction to the models used here, see Hyvärinen et al. (2009).

<sup>11</sup>The theory of unsupervised learning is much less developed and more complicated than the theory of supervised learning, especially for multi-layer networks. Therefore, a lot of work on such feature learning uses supervised learning, somehow obtaining labels or categories for each image, and using ordinary supervised learning where the network learns the connection between the images and their categories. The bottleneck here is getting sufficient amounts of such data with category labels. It is difficult because somebody has to tell what the photographs are depicting; if the labels are given by humans, that is a lot of work (although a simple approximation

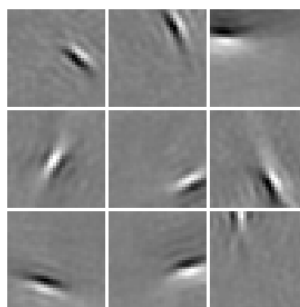


Figure 10.3: Simple image features learned by a neural network. Each small patch gives the synaptic weights in a neuron whose input consists of small image patches. The weights can thus be plotted as grey-scale values arranged as image patches. More precisely, these are the results of applying a method of unsupervised learning called independent component analysis on small image patches.

prior information about images. In general, the multi-layer network will be computing increasingly complex features in each layer.<sup>12</sup> The features computed by the units in higher layers are no longer simple edges or bars: they are more like some specific parts of the objects that the network was trained on. They are also more focused on coding the identity of those parts while ignoring less relevant details such as wherein the image the parts are located. For example, a neuron in a high layer could respond to a cat head, irrespective of where it is in the input image, and further ignore details such as the exact shape of the face of the cat. In this sense, such neurons are quite similar to cells in the inferotemporal cortex, an area in the brain that performs a very high level of image analysis.<sup>13</sup>

would be to extract the labels from captions which are sometimes attached to images on the internet). Current research is strongly focused on finding methods to train multi-layer neural networks without labels, that is, in an unsupervised way. A particularly promising approach is called *self-supervised*, which means performing unsupervised learning by reformulating the problem as supervised learning. Basically, you create hypothetical outputs, or a hypothetical classification problem, and use them to train your ordinary supervised, input-output neural network. The possibilities are unlimited: you could define the input to a neural network to be a degraded version of your data and the output your real data, where the degraded version could be obtained by adding noise, or making a colour image black-and-white (Vincent et al., 2008; Larsson et al., 2017). Or, the “degraded” data could actually be artificially generated: then you train the neural network to distinguish between the real and the artificial data (Gutmann and Hyvärinen, 2012). For example, in video data, you could randomly shuffle the time frames in a video, or scramble audio in a video with sound, and train the neural network to classify such scrambled data vs. the original data (Hyvärinen and Morioka, 2017; Misra et al., 2016; Arandjelovic and Zisserman, 2017). In each case, the neural network has to learn something about the structure of the data in order to perform this mapping, that is, trying to reconstruct the original images from degraded ones, or telling which data is real and which is noise. The multi-layer processing thus learned are reasonably similar to what is computed in the brain (Zhuang et al., 2021). However, it should be noted that self-supervised learning in itself gives only features; it does not give a proper Bayesian prior model except in some special cases, such as the “noise-contrastive estimation” by Gutmann and Hyvärinen (2012), and nonlinear versions of independent component analysis (Hyvärinen and Morioka, 2016; Khemakhem et al., 2020).

<sup>12</sup>(Güçlü and van Gerven, 2015; Kriegeskorte, 2015; Eickenberg et al., 2017; Zhuang et al., 2021)

<sup>13</sup>(Tanaka, 1996; Brincat and Connor, 2004). The property of ignoring less meaningful details is called *invariance*. There have also been claims that neurons in the inferotemporal cortex, which performs a very high level of image analysis, could be coding for the identities of single people (Quiroga et al., 2005). However, a more detailed analysis of the results shows that this is an exaggerated interpretation: single neurons are probably responding to several different people (Quiroga et al., 2008).

### Illusory contours

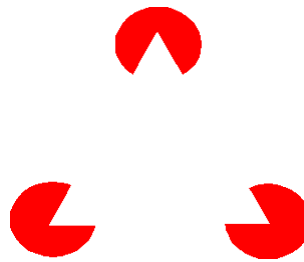


Figure 10.4: The Kanizsa triangle, an example of a visual illusion. There is actually no triangle in the figure, just pacmans.

### Illusions as inference that goes wrong

We have now seen that the incompleteness of the incoming sensory information can be, to some extent, alleviated by Bayesian inference. However, this solution is far from perfect—whether we consider perception in humans or sophisticated AI. Sometimes the perception is blatantly incorrect, as shown by the phenomenon of visual (or “optical”) illusions. A dramatic example is shown in Figure 10.4. We tend to see a full triangle in the figure, with three uninterrupted lines as its sides or edges. In reality, though, the sides of the triangle do not exist in the figure. If you cover the “pacmans” with your fingers, you see that there is nothing but white space between them. Yet, most people have a vivid perception of three lines between the pacmans which create a complete triangle.

This is called an illusion in neuroscience since the sides do not physically exist in the figure; they are simply imagined by our visual apparatus. Just like the imagination of a full disk in Figure 10.2 we saw earlier, this can be considered unconscious inference, where your visual system computes the most likely interpretation. The difference is that here, the interpretation is in clear contradiction with the actual stimulus, or physical reality. While inferring a full disk in Figure 10.2 seemed smart and would quite probably have been correct in real life, inferring that there is a full triangle in Figure 10.4 may seem quite stupid, at least after you have checked that the sides do not really exist. The curious fact is that you cannot really help seeing the triangle in Figure 10.4.

The theories explained in earlier chapters help us further understand why such illusions occur. A neural network is trained to accomplish a well-defined task, such as recognizing different objects in photographs. However, such neural networks are inflexible and only able to solve the problem they are trained for; neural networks are not general problem-solving machines. In particular, a neural network will not work very well when the input comes from a different source than what it was trained for. Arguably, the Kanizsa triangle is something artificial, and different from what you would usually see in real life (where pacmans are quite rare), so we should not expect the brain’s neural networks to process it appropriately. This is another way of saying that the brain’s prior information contains assumptions that are typically true in the context where you usually live, but they are just about probabilities, and might sometimes turn out to be quite wrong.

At the same time, the dual-system theory explains why it does not help if somebody explains to you that this is an illusion, or even if you realize that yourself. A logical, symbol-level understanding that there is no triangle has little effect on the other system, i.e., neural networks, which are mainly in charge of visual perception.

## Attention as input selection

As already mentioned, in real life, any sophisticated perceptual system further faces the problem that there is simply too much information in the visual field. This problem is very different from the missing information problem solved by the prior information, as just discussed. In particular, there are often too many things in the visual input at the same time. There may be many faces, people, buildings, animals, or cars, at the same time, and it is too difficult to process all of them. This is in stark contrast to current success stories of object recognition by AI, which are usually obtained in a setting where each input image contains only one object, or at least one object is much more prominent than the others. Suppose you input an image of a busy street to such a neural network trained to recognize a single object in an image. Since the input now contains many objects, features of different kinds will be activated in the neural network, some related to the perception of people, some related to the perception of buildings, some to cars, and so on. Many of the features are actually quite similar in different objects: think about two faces in a crowd, which are quite similar on the level of pixels and even rather sophisticated features. Various neurons will be activated, but it is impossible to tell which were activated by which face. It will be very difficult for the AI to make sense of such input and the activations of its feature detectors.

This problem really arises when the information processing as well as the input data sensors work in a parallel and distributed mode. Parallel and distributed processing, considered in detail in Chapter 11, usually means that there are many processors which work simultaneously and independently. Here, the situation is even more extreme since the input data itself is received from a huge number of sensors, such as pixels in a camera or cells in the retina. Yet, the principles of parallel and distributed processing are really the same, as the outputs of such sensors are processed by a large number of small processors, at least in the brain.

Traditional computer science usually does not deal with this problem. If the input to the computer is mouse clicks by a human user, the input is quite manageable. Even if a computer handles a very large database, the situation is different because it follows explicit instructions on what information to retrieve and in what order. Vision is more like thousands of disk drives simultaneously and forcefully feeding the contents of their databases to a single computer.

The key to how the brain solves this problem, especially in the case of vision, is the multi-faceted phenomenon of *attention*. In the most basic case, the visual system of many animals, including humans, selects just one part of the input for further processing. As we say in everyday English, the animal only “pays attention” to one object at a time, whether it is a face seen on the street, or some object it is trying to manipulate.

The simplest form of such *selective* attention is that you just wipe out everything else in the visual field, except for one object. In Figure 10.5, we see a photograph and an attentional selection template, which shows how only the main object of interest in the figure is found and selected. The results of such computation can be used to simply blank out everything else except for the main object.<sup>14</sup> Such a form of attentional selection is also called “segmentation”. Now, if you input an image that contains only this one object to the neural network, the recognition will be much easier. The most amazing kind of attentional selection that our brains can accomplish must be finding individual faces in a crowd. Face processing is evolutionarily extremely important, so there are specialized areas in the human brain for processing just faces (monkeys have them too).

Performing such segmentation is not easy: using attentional mechanisms in AI and robots is an emerging

---

<sup>14</sup>(Borji et al., 2015; Zhou et al., 2019; Chen et al., 2018)

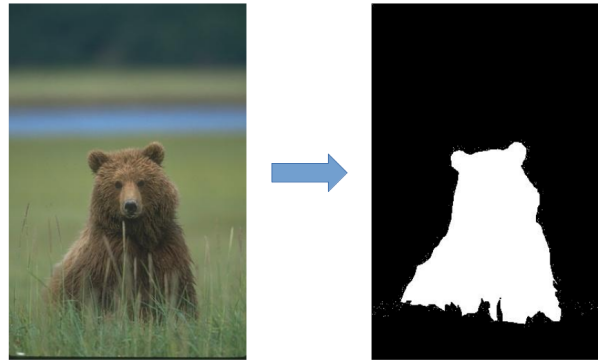


Figure 10.5: Attentional selection illustrated. The photo on the left is the original visual input. An attentional system selects the pixels to retain, shown in white in the figure on the right. (Based on data by Martin et al. (2001), used with permission.)

topic, and we still don't know very well how to do it. However, like so many other functions related to intelligence, it may be possible to learn it.<sup>15</sup> In fact, such selection seems to be happening in many different parts of the brain and in many different ways. In a sense, it is a reflection of the ubiquity of parallel processing in the brain, which necessitates various forms of input selection all over the brain.<sup>16</sup>

The key point in attention is that it leads to a bottleneck in the processing. In the example just given, only the one single object left in the image is given to further processing, including the final pattern recognition system. So, only one object can be recognized at a time, since all the others are wiped out. It is often said in cognitive neuroscience that “attentional resources are limited”, and here we see one illustration of that principle: if you pay attention to one thing, you will necessarily tend to ignore everything else. This, in its turn, increases uncertainty since you don't know much about those things you are not paying attention to.

## Subjectivity and context-dependence of perception

An important consequence of the uncertainty of perception is its subjectivity: I see one thing, and you may see something different. Being based on unconscious inference using prior information, perception is subjective if different people or agents have different priors. Then they will interpret the incomplete incoming information in different ways.

The priors used in human perception actually contain many different parts. There is one rather permanent

<sup>15</sup>(Minut and Mahadevan, 2001; Mnih et al., 2014; Greff et al., 2016). Attention is fundamentally a form of action: even moving your eyes can be seen as a form of attention, since it helps to select certain parts of your environment for visual processing. Thus, learning to attend may be possible by the same principles by which an agent can learn to act in intelligent ways, discussed in Chapter 3.

<sup>16</sup>The word “attention” is quite overloaded with different meanings in cognitive psychology. The sensory attention we have considered here is very different from some other kinds of attention. In particular, another type of attention very relevant for this book is *sustained attention*, considered in Chapter 9, which means you try to concentrate on a single task, such as reading a book, for an extended period of time. That is very different from sensory *selective* attention considered here since sustained attention is about *long-term* attention on a *task* instead of relatively short-term attention on sensory objects. Selective attention can further be divided on another axis: bottom-up attention, where an external stimulus grabs your attention (as in the case of interrupts in Chapter 8), and top-down attention, used for example when you search for a certain person in a big room and only pay attention to faces. (The exact terms used in the different cases are quite variable in the literature.)



and universal component, shared by all humans, and probably many animals. It includes those general regularities that can typically be found by training artificial neural networks. But another component in the prior is more individual and depends on the experience of the agent (animal, human or AI). When an agent observes things happening, ideally it will incorporate all the new observations into its prior—possibly after performing some kind of attentional selection. If it didn't, it would be wasting valuable data that it has collected on the world. It is this individual part of the prior, based on their own experiences, that makes the priors different from one agent to another. Each agent may even be living in a different environment; they may spend their time in very different occupations. So, it is clearly useful that the prior is different from one agent to another. But this necessarily implies that perception will be different as well. You don't see exactly the same thing as your friends, not to even mention your robot. This might not be such a serious problem if the agent understood the subjectivity of perception well enough. However, such understanding often escapes even humans.

There is even a further component in the prior, which depends on the context, e.g., where the agent is at the moment of perception. If you're at home, you expect to see certain kinds of things, and if you're walking on the street, you expect to see other kinds of things. This leads to dependence of perception on the context, even for the same agent.<sup>17</sup> These limitations of perception reflect the limitations of categories discussed in Chapter 7. Categorization is usually based on perception, so if perception is subjective and context-dependent, the categories inherit those properties as well.

Perception is made even more subjective by the selection of incoming information by attentional mechanisms. Attention has a huge impact not only on the immediate perception in the agent, but also on the model it learns on the world. Fundamentally, attentional mechanisms choose the data that is input into the learning system. Anything not attended is pretty much ignored and not used in learning. As our brain “creates our world” in the sense of reconstructing it from sensory input, that creation is thus strongly influenced by attentional mechanisms.

## Reward loss as mere percept

A crucial insight that this view on perception gives to suffering is that its causes are subjective and generally uncertain: they may be based on faulty inference. In particular, reward loss is just another percept (i.e. a result of the process of perception). It is based on solving an inverse problem using prior information to compute a percept of the obtained reward.

Misperceptions of rewards may be particularly common when perception of other people's reactions are involved. You might perceive the facial expression of your friend as angry, and register some negative reward as resulting from your actions. But perhaps your friend just had a bad headache and his face reflected that; taking the uncertainty of perception into account should help you behave in a more appropriate way towards him.<sup>18</sup> An extreme example of misperception of reward is found with some drugs of abuse. They feel good and you perceive a reward on a biological level. Yet, such perception has no real basis: The drug merely misleads your

<sup>17</sup>(Bar, 2004). Even the perception of pain is modulated by context and history (Tabor et al., 2017); see also Chapter 6 and its footnote 19. Perception can also be modulated by metabolic states, such as hunger (Livneh et al., 2017). There are also claims that desire and aversion (motivational states) could directly influence perception (Balcetis and Dunning, 2006), but such phenomena are controversial (Firestone and Scholl, 2014).

<sup>18</sup>It is also typical for people to value objects more if it takes a lot of effort to obtain or produce them. This can be seen as a simple heuristic to approximate the reward, but it can of course go wrong (Kruger et al., 2004).

brain into perceiving a reward by perturbing its metabolism.<sup>19</sup>

The situation is even more complicated here since the agent also computes the expected reward based on the information it has at its disposal and using the available computational capacities. Thus, there seem to be two different ways in which uncertainty in perception affects the computation of reward loss: the obtained reward may be perceived wrong, or the computation of expected reward may go wrong.<sup>20</sup> Both are just logical consequences of computation performed with limited resources and limited data. Ultimately, a reward loss may even be illusory in the sense that one is perceived but it is merely a mental construct. One way to remedy the situation is that we should actually consider *perceived reward loss* instead of any objectively defined reward loss, since the agent can never know with certainty what the reward loss was. (We will postpone the details of such a re-definition to Chapter 13.)

## Ancient philosophers on perception

The uncertainty and subjectivity of perception were discussed by several ancient philosophers. In ancient Greece, the Skeptic school was particularly prominent in pointing out the limits of human knowledge, including the relativity of perception. The Pyrrhonian branch was fond of giving examples where different people perceive the same thing differently:<sup>21</sup>

When we press the eye from the side, the forms and shapes and sizes of the objects we see appear elongated and narrow.

The colour of our skin is seen as different in warm air and in cold, and we cannot say what our colour is like in its nature, but only what it is like as observed together with each of these [circumstances].

Such uncertainty leads the skeptic to adopt an attitude of not making any judgements on external objects:

So, since so much anomaly has been shown in objects (...), we shall not be able to say what each existing object is like in its nature, but only how it appears (...) therefore, it is necessary for us to suspend judgement on the nature of external existing objects.

A Japanese Yogācāra-inspired poem beautifully describes a scene where different agents have very different interpretations of the same sensory input:<sup>22</sup>

At the clapping of hands,  
the carp come swimming for food;  
The birds fly away in fright, and  
A maiden comes carrying tea—  
Sarusawa Pond

---

<sup>19</sup>(NIDA, 2020)

<sup>20</sup>However, see footnote 17 in Chapter 14 on whether it makes sense to say that expectation of reward is “wrong”.

<sup>21</sup>Sextus Empiricus’s *Outlines of Pyrrhonism* from ca. 200 CE, with translation taken from Annas and Barnes (1985), see also e.g. (Morison, 2019; Mates, 1996).

<sup>22</sup>(Tagawa, 2009)

When somebody claps his hands by the famous Sarusawa Pond in Nara, Japan, the carps interpret it as a call for feeding; the birds are scared of the noise and flee; while a maid of a near-by inn thinks a customer is calling for her.

It is perhaps easy to admit that an animal or a robot sees things differently from yourself, either in a more primitive way, or in a superhuman way. Yet, it is notoriously difficult for humans to admit that two people can see the same thing in different ways, and that both ways can be equally valid. But there is something even more challenging; there is an even more difficult implication of the theories discussed in this chapter. It is the general idea that all our perceptions are actually just interpretations, or beliefs, or inferences, instead of revealing an objective truth. In AI theory, it is never claimed that the agent *knows* anything; the very concept of knowing is conspicuously absent in that theory. All an AI agent has is *beliefs*, and those are usually expressed in terms of probabilities, lacking any certainty.

When you take this line of thinking further, you may arrive at the idea that *all* we believe or pretend to “know” is based on our perceptions, and thus inherits the uncertainty and the subjectivity of perception. In fact, one could say that my perception defines my world. This may actually be rather obvious to anybody who programs a sensory system in an AI. Such ideas are often associated with Asian philosophical systems such as Mahayana Buddhism, especially the Yogācāra school and later schools drawing on those ideas, including Zen.<sup>23</sup> Yet, those ideas have also been beautifully expressed in the West, where their foremost proponent may have been David Hume who wrote:<sup>24</sup>

Let us chase our imagination to the heavens, or to the utmost limits of the universe; we never really advance a step beyond ourselves, nor can conceive any kind of existence, but those perceptions, which have appeared in that narrow compass. This is the universe of the imagination, nor have we any idea but what is there produced.

We will see these deep points re-iterated and expanded in later chapters, especially Chapter 12.

---

<sup>23</sup> (Williams, 2008b)

<sup>24</sup> (Hume, 1739), Section 1.2.6. In fact, Gopnik (2009) argues that Hume’s ideas may have been influenced by Buddhism through some Jesuits; see also footnote 34 in Chapter 7.

## Chapter 11

# Distributed processing and no-self philosophy

The concept of a “self” is central for understanding suffering, but it is highly complex. Some aspects of self were already considered in Chapter 6. In this chapter, I consider another central aspect of self, related to control. Self can be seen as the entity that is in control of actions, including control of cognitive operations inside the agent, or, to put it simply, in control of the mind.

In preceding chapters, we have seen cases where the mind seems to be difficult to control, due to automated interrupts and wandering thoughts. Here, I consider a general cognitive principle that explains why control is limited. The idea is that in the case where the information processing is parallel and distributed, it is difficult for any single part of the agent’s information-processing system to be in charge of the whole system, e.g. the whole brain. This massively parallel and distributed nature of the brain thus creates most of the uncontrollability in the human mind. The lack of control considered here can also be seen as a generalization of dual-process nature of the mind considered in earlier chapters. Here, there are not just two processes competing for control, but a great number of them.

These considerations necessarily lead to the question of free will: can an AI, or even a human, actually have free will—and what does that mean in the first place. From the viewpoint of the theories of perception in the preceding chapter, we can ask if perception of control and free will are simply illusory perceptions, thus providing another link between the uncertainty of perception and uncontrollability. Such considerations have lead some philosophers to propose that there is no self, or no doer of actions, and I will revisit these ideas from a computational viewpoint.

### Are you really in control?

Suppose you just raise your arm—you can physically do it while reading this if you like. You probably think it was you who decided to raise the arm, and it was you who actually executed the action. You felt being able to control the world, or at least your arm in this case.<sup>1</sup> This “you” that first controlled your mind by making a decision, and then controlled your arm, is what can be called the self—in one meaning of the word. The self chooses actions, and controls some aspects of the world, including your inner world.<sup>2</sup>

---

<sup>1</sup> Philosophers talk about (the feeling of) “agency” (Metzinger, 2003) . I don’t use that terminology because it would lead to confusion in this book where the word “agent” usually means something different. Furthermore, such agency is usually related to a conscious feeling, while in this chapter, I refrain from talking about anything related consciousness, which will be treated separately in Chapter 12.

<sup>2</sup> (Skinner, 1996)

However, a number of thinkers have proposed that in fact, “you” are not really in control of anything. A case in point is wandering thoughts. It can be claimed—following a strict definition of the term—that we never want to have wandering thoughts: if we want to think what we are actually thinking, the thoughts are not called wandering. Furthermore, wandering thoughts often feel unpleasant, for example in the extreme case of rumination. So, why do we then continue having them?

A well-known experiment on the control of thoughts is to try to *not* think of a pink elephant. This is another exercise you can do right now: for a minute or so, do not think of a pink elephant. What invariably happens is that you will be thinking of a pink elephant in spite of your trying not to, or perhaps precisely because of that trying. Clearly, our control of thoughts is limited. In addition, interrupts such as fear, anger or desire capture our mind and direct the processing in ways we might not want. Even habitual behaviour can be seen as a lack of control in some cases: if you mindlessly follow habits, you may end up doing something you would not have done if you had actually deliberately planned your actions.

Lack of control increases suffering in our basic framework of suffering as frustration. Lack of control reduces the probability that the agent reaches the goals it has committed to; it cannot get the things it wants, or avoid the things it is averse to. That means there will be more frustration and reward loss. In fact, the very existence of suffering can be seen as a form of uncontrollability, since if you could really control your mind, you would probably just switch off any feelings of suffering.

### *Philosophical views on uncontrollability*

In philosophy, the idea of lack of control and its connection to the self goes back to, at least, the Buddha’s times. In a famous discourse, he explained why there actually is no such thing as “self”. He started his refutation by considering the human body, saying<sup>3</sup>

[I]f the body were self, the core of our being, then it would not tend to affliction or distress, and one should be able to say of it, ‘Let my body be thus (in the best of conditions); let my body not be thus (in a bad condition).’ It should be possible to influence the body in this manner.

He continued by going through different aspects of the human mind (perception, thinking, etc.), and denying that any of them could be called the self either, since none of them can be properly controlled. For example, “no one can wish for and manage thus: ‘Let my perceptions be thus, let my perceptions be not thus’”. If you smell something disgusting, you cannot just decide not to smell it.

Thus, originally, the Buddha framed the very concept of self in terms of control: self is what is in control.<sup>4</sup> Since, as he argues, there is actually no (or little) possibility of control, there can be no self. Realizing this is thought to be essential to reduce suffering.<sup>5</sup>

In ancient Greece and Rome, the Stoic philosophers had similar ideas. Perhaps the very core of Epictetus’s philosophy is contained in his attitude towards control:<sup>6</sup>

<sup>3</sup>(Mahasi, 1996), based on *Samyutta Nikaya* 22.59; with explanatory text in parentheses added by Mahasi.

<sup>4</sup>Arguably, the Buddha’s viewpoint could also be interpreted as the self being what can *be controlled* instead of what *controls*. Nevertheless, according to Mahasi (1996, p. 12–14), what the Buddha is denying is precisely a “controlling self” as well as an “active agent self”. The two viewpoints are in a sense unified when Harvey (2009, p. 49) proposes that according to Theravadan Buddhist thinking, “a Self would have total control over itself.” In any case, this makes little difference in what follows where the main point is a general lack of control, or uncontrollability.

<sup>5</sup>(Verhaeghen, 2017; Harvey, 2009)

<sup>6</sup>The very first lines in *The Enchiridion*.

Some things are in our control and others not. Things in our control are opinion, pursuit, desire, aversion, and, in a word, whatever are our own actions. Things not in our control are body, property, reputation, command, and, in one word, whatever are not our own actions.

Epictetus's idea of uncontrollability is more limited: we cannot control what others do or think about us, or, in line with the Buddha, our bodies. But in stark contrast to the Buddha, he thinks we can control our thoughts and feelings, including desires and aversion. Presumably, Epictetus did not practice the same kind of meditation as the Buddha, which might have convinced him of the uncontrollability of thoughts and feelings. In any case, both philosophers advocated recognizing how little control we have as a means of reducing suffering—we will discuss such practical implications in Chapter 14.

We seem to actually have two different kinds of uncontrollability here. First, the uncontrollability of the outside world as emphasized by Epictetus; and second, the uncontrollability of the mind as emphasized by the Buddha. The uncontrollability of the outside world is easy to understand, and its causes are rather obvious. The agent has limited strength: it probably cannot lift a mountain. It has limited locomotion: if it is designed to move on wheels, it probably cannot fly. If it lives in a society, it has limited means of influencing other agents.

What is less obvious, and my focus here, is that there seems to be so much uncontrollability regarding the mind. We have already seen examples where control of the mind is lacking, as in the case of interrupts and wandering thoughts; the dual-process structure of the mind creates further conflicts and reduces control. Therefore, the question arises whether there is some *general principle* behind all of those manifestations of uncontrollability.

## Necessity of parallel and distributed processing

The basic idea here is that the lack of control of the human mind is fundamentally based on one property of the brain: parallel and distributed processing. That is, there are many processors, or neurons, processing the information at the same time, and to some extent independently of each other. If there are many processors working independently, each of them cannot be in control of the agent's actions: there has to be some kind of arbitration, at the very least. Modern AI also uses such parallel and distributed processing, in particular in the form of neural networks. Both the brain and neural networks in AI are in this way fundamentally different from an ordinary computer, which typically uses serial processing in a single processor.<sup>7</sup>

While these properties have been mentioned in earlier chapters, we have not really considered the question of *why* parallel and distributed processing happens. From a biological perspective, we need to find some evolutionary justification for why the brain is parallel and distributed, and from a computer design perspective, we need to explain why such processing would be useful. Perhaps we can answer both questions if we simply find some fundamental computational advantage in parallel or distributed computation.

### *Failure of Moore's law and necessity of parallelization*

Let's first consider the question of parallel processing from an AI viewpoint: What is the point in using many processors? If you want to speed up your computations, why not just get a single processor which is, say, a hun-

---

<sup>7</sup>In practice, a personal computer or a mobile phone would not usually have just one single processor, but a small number of them, typically less than ten. For example, the display would be supported by a separate processor, a graphics processing unit. Merely for the purpose of keeping the discussion simple, I will assume an ordinary computer has just a single processor.

dred times faster, instead of putting together one hundred more ordinary processors that compute in parallel? Obviously, there is a limit to how fast processors you can buy for an AI. Perhaps you need faster computation than what is given by the fastest single processor available today. That is why all the supercomputers in the world are highly parallel; they are collections of thousands of processors. That is the only way to increase the computational power to record-breaking extremes.

On the other hand, if you're really lazy, you might be tempted just to wait. We all know that the technology behind the processors has been developing at an enormous speed. The famous Moore's law states that the computing power of a processor doubles every two years. This may lead to the impression that there is really not that much reason to go through the trouble of parallelization: if the fastest processor is not fast enough, just wait a few years, and it will be. If this logic were true, it would also mean that there may not be any fundamental reason why computation in AI needs to be parallel, since the power of a single processor seems to grow exponentially and without limit.

Yet, there are fundamental reasons why really efficient computation may not be possible at all without parallel computation, and why, in fact, Moore's law is not true anymore. One reason is that making processors faster is to a large extent driven by making them smaller. A smaller processor means shorter delays in transmitting the information inside the processor. Such miniaturization cannot go on forever because at some point, you get too close to the level of single atoms, and the laws of physics basically change in the sense that quantum phenomena start appearing.<sup>8</sup>

A more practical problem is that due to complicated physical phenomena, faster single processors use much more energy than a set of slower processors with the same total computational capacity.<sup>9</sup> Energy is obviously expensive and cannot be used in unlimited quantities. Moreover, such an increase in energy consumption has another, surprising effect, which is that the processors heat up very quickly, and keeping processors cool is increasingly becoming a problem. If you design a new processor which is ten times faster than your current one, the power consumption and the heat generated are usually much more than ten times larger.

So, these are convincing reasons why it is necessary in AI to use many processors in parallel. In fact, the speed of a single processor ("clock rate") even in mainstream computers has not been really increasing since around 2005. The overheating problem became so serious that faster processors became impractical to use.<sup>10</sup> Since splitting the computations into many processors generates less heat, manufacturers started putting together several processors on a single chip—the processors are now called "cores". The number of cores in an ordinary computer is still small, though, so this is very far from the massively parallel case seen in the brain.

### *Parallelization can be hard*

Thus, the great promise of parallel processing is that it can be much faster than serial processing, given the same budget of energy, or, indeed, money. However, there is a problem. If you have one hundred processors that process the same information at the same time, the processing could be, in principle, a hundred times faster. But that only happens in an ideal scenario which requires that the computations are such that they can be parallelized, i.e. they can be simultaneously performed on one hundred processors without any problems. Some problems can easily be parallelized, while others are more difficult, perhaps even impossible. Program-

---

<sup>8</sup>(Theis and Wong, 2017)

<sup>9</sup>(Markov, 2014)

<sup>10</sup>(Markov, 2014; Gorder, 2007)

ming parallel systems needs special algorithms, as well as specialized expertise.

Consider a problem of finding a small object, say a single very black pixel, in an input image. (Suppose for simplicity there is only one such object in the image). You could have a single serial processor scanning the image pixel by pixel. That might take, say, 100 microseconds (one microsecond being one-millionth of a second). On the other hand, if you have 100 processors, you could split up the image into 100 regions, and tell each processor to search for the pixel in one of the regions, and then report to a central processor whether it was there or not. That should not take much more than 1 microsecond. This problem is easy to parallelize, and the speed-up (100x) is basically the same as the factor by which you multiplied the number of processors (100x). A neural network, whether in AI or in the brain, can do such computations massively in parallel, and thus incredibly fast. This is one of the reasons for the impressive behaviour of the human visual system, and the success of neural networks computer vision tasks.<sup>11</sup>

Then there are tasks that are really difficult to parallelize. This is generally the case when you need to compute an intermediate result before proceeding further. As an intuitive example, consider building a house with rather traditional methods. You first have to build a foundation, and let it dry. Then you build the walls, and finally, set the roof. Suppose you had an unlimited number of builders that you can use; telling them what to do is like trying to parallelize computation. Now, the problem is that you cannot meaningfully divide the builders into three teams so that one of them sets the roof at the same time as another group lays the foundation! Also, if you really have a huge number of builders, they would not even fit on the building site. So, parallelization can be tricky.

Optimization by a gradient method is an example of something that is typically considered difficult to parallelize because you need to do it step by step. Yet, a lot of effort has been spent in computer science research to figure out methods that enable parallelization of such algorithms, sometimes quite successfully.<sup>12</sup> With a lot of intellectual effort and ingenuity, it is possible to parallelize even seemingly impossible problems. However, such parallel methods are quite complicated, and thus parallel computation in gradient methods is not currently very often.

The fact that some computational problems are hard to do in parallel while others can be parallelized very efficiently is part of the reason why ordinary computers and the brain are good in very different things. The brain is particularly good at vision, for example. Vision can be rather easily parallelized, as was seen in the simple pixel-finding example above, and indeed the best AI solutions to vision have imitated the brain using neural networks. On the other hand, ordinary computers are very good in logic-symbolic processing, as discussed earlier.

But what is the evolutionary import of these considerations—does it make sense to claim that the brain is massively parallel because of the above-mentioned reasons related to the clock-speed of processors? Certainly, the constraints in building an intelligent system with biological hardware are very different, and the logic above may be mainly relevant for AI. What it actually shows is that progress in AI seems to need computers which are more and more similar to the brain. Yet, it is possible that the massive parallelization in the brain might have

---

<sup>11</sup>To take another example: Tree search, which is essential planning, can also be parallelized, although it is a bit more difficult. After a few steps in the search tree, you can distribute the different branches to different processors, and each processor can search further in one of the branches. Intuitively, this would be like a boss assigning five different scenarios to her employees in a planning exercise. Each employee gets one scenario which each start from different assumptions, corresponding to the first branchings of the search tree. After the initial hurdle of formulating the scenarios (that is, building the initial branches of the tree), the parallelization is easy.

<sup>12</sup>(Zinkevich et al., 2010; Recht et al., 2011)



some relation to the energy-efficiency considerations that we just saw.

### *Distributed processing reduces need for communication*

The second question is why distributed processing is needed. Distributed processing is different from parallel processing in that the emphasis is on different processors working independently with as little communication as possible. Distributed computing is important, even necessary, simply because communication is often quite expensive. In the brain, most of the volume actually consists of white matter, which is nothing else than “wires” (called “axons”) connecting different neurons. Those wires take up much more space than the actual processing units. So, the sheer space available in the head strongly limits the connectivity of brain neurons.<sup>13</sup> In addition, communication consumes energy which is, again, another limiting factor.

What makes achieving full connectivity particularly difficult is that the number of possible connections between processing units grows quadratically as the number of processing units grows. If you have a million processors, and you want to build connections between all the possible pairs, you need almost a trillion (1,000,000,000,000) wires (assuming each wire can transmit information in one direction only, as happens in the brain). So, the amount of connections easily becomes a limiting factor, and it is important to do the computations needed using minimal information transfer between the processing units, by judiciously designing the way the different areas are connected with each other.<sup>14</sup>

This is the central point about distributed processing: When communication between the processors is expensive, special solutions are needed. In AI, there is a thrust to distribute AI computation to smartphones that collect the data in the first place, so that the amount of data they transmit to each other or any central server would be minimized.<sup>15</sup> In the brain, part of the solution is that processing is very clearly distributed on the level of large brain areas. There are areas responsible for processing visual input, areas for processing auditory input, areas responsible for moving the muscles, areas for spatial navigation, and so on. Each of these areas does its computations relatively independently. That is possible partly because they get different input (visual vs. auditory), and partly because they need to solve different computational tasks (object recognition vs. moving muscles). The communication *between* those areas can then be strongly limited, and less wiring is needed.

Distributed processing will create its design problems, just like parallel processing. Some tasks are easy to distribute over processors, while others are less so. Again, neural networks are an example of processing which is highly, even massively distributed, and clearly works well in applications such as sensory processing of images and sounds. Considering the example of finding a small object in an image described above, it is easy to see that the computation described is also strongly distributed since the 100 processors each get their own input and then do their computations with no communication between them needed.

## **Central executive and society of mind**

The logic above suggests that sophisticated intelligent agents may have to be a collection of relatively independent parts or processors—and that is certainly the case in the brain. The resulting computing system is very

---

<sup>13</sup> (Zhang and Sejnowski, 2000; Hari, 2017)

<sup>14</sup> (Bullmore and Sporns, 2012)

<sup>15</sup> (Xu et al., 2020)

different from the view we intuitively have of ourselves. We tend to think of ourselves as serial processors because much of our inner speech and conscious thinking is serial. Speech is inherently serial because the words follow one after another in one single “train” of thought. But such introspection, based solely on what we can consciously perceive, is quite misleading.

A simple metaphor for illustrating the counterintuitive properties of a parallel and distributed system is the “society of mind”: the different mental faculties are compared to human individuals that together constitute a society which is precisely the mind.<sup>16</sup> One individual (or processor) is monitoring, say, the state of the bowels, and another one is, independently, responsible for recognizing the identities of faces whose images are transmitted by the eye. Those processors are like human workers with well-defined, separate tasks. Each one may be active much of the time, thus working in parallel. In line with the computational arguments we just discussed, it may also be intuitively clear that it is important that the different individuals mind their own business most of the time, focusing on their own part of the work. Therefore, they only interact if it is really necessary, with minimum communication; thus, the operation is distributed. This metaphor is trying to counteract the intuitive impression we tend to have that the mind is a single, serially processing entity which would be difficult to divide into parts.

Now, to see the point regarding control, consider whether it is possible that one of the independent processors is actually in control of all the others. Psychological theories often use the term *central executive* for that part of the mind which is supposedly in charge, controlling the rest.<sup>17</sup> At first sight, having such a central executive sounds like common sense. The brain has many sensory processing systems (vision, audition, etc.), it can send commands to a multitude of muscles to execute actions, and above all, it has complex information-processing capacities in terms of planning and learning. It would seem that such a system must fall into complete chaos unless there is one area which controls the others. That would be the central executive, a brain area that controls all, or at least most, of the other areas. It would integrate information coming from them and, in return, send processed information and commands to each of them. In the society of mind metaphor, this would correspond to a leader of the society that tells all the individuals what they should do.

It could be argued that having a single area to control all the others is to some extent in contradiction with the whole point of distributed and parallel processing. The central executive would need to have particularly great processing power, and it would need to receive a huge amount of information from all the other parts of the whole system. Thus, both the two bottlenecks discussed above, processing speed and communication capacity, would resurface—but we will see below that this is not really the case.

Designing such a system with a central executive is not very different from designing different decision-making systems in a human society or organization. If there is a single leader, she must inevitably delegate a lot of power to others (say, ministers) in order to reduce the processing power needed by herself. Then, the leader is strongly dependent on the information passed on by the ministers; the leader does not have enough time to make decisions on all the details. So, the power of the central executive is limited due to the limitations on the computational power of a single processor.

On the other hand, if there were a central executive, what about wandering thoughts, emotional interrupts, or habits? Is the central executive just watching when the whole system is hijacked by the fear elicited by the sight of, say, a spider? We argued in earlier chapters that emotional interrupts are useful for evolutionary

<sup>16</sup>This is a rather liberal interpretation of Marvin Minsky’s original idea (Minsky, 1988; Singh, 2012). For Minsky, the individuals (which he calls “agents”) are very simple, more like subroutines in a computer program, as opposed to humans.

<sup>17</sup>(Baddeley, 1996)

purposes, so the leader might actually not be very unhappy about that. But interrupts, by their very nature, cannot be prevented, not even by the central executive. Is there any point in calling such a leader the central executive if she is not really controlling everything that happens? What if you eat chocolate because you have a habit of doing it every day (in addition to an irresistible desire, perhaps), even though one part of you knows it is bad for you in the long run—who actually made that decision?

This logic has led many to the proposal that in the human mind and brain, there is no central executive, or, metaphorically speaking, *the society of mind has no leader*. That is, there is no part in the mind that controls the rest, nothing that controls everything else that happens in the society.<sup>18</sup> The society is fundamentally a collection of relatively independent actors. This means very concretely that there is no particular part of the mind or the brain that would control our thoughts, feelings, or desires: they just come and go depending on a complex interaction between different brain areas. Each part of the mind can propose its own mental actions. One part of the visual system might tell the motor cortex: “Let’s move the eye gaze to the right since there seems to be something interesting there”, but at the same time, the replay system might insist on replaying a past episode while ignoring whatever may be happening in the outside world. The result may be a bit chaotic, and having, say, wandering thoughts would not be surprising. To the extent that we define the self as the central executive, there would be no self, in line with Buddhist philosophy.

While such a philosophy is fascinating, it has to be pointed out that there are also neuroscience results claiming that some brain regions in the prefrontal cortex are actually the central executive.<sup>19</sup> Moreover, in the design of distributed computing architectures, it is well-known that having some kind of a central processor actually makes communication easier. The point is that there is a good compromise to be found between the two extremes of completely distributed computation and computation in a single processor. Such a compromise can in fact be found in computation which is mainly parallel and distributed, but, crucially, includes a central processor that *coordinates* the computation, which is still mainly performed by the other processors. In the example above, with a million processors, we saw that a fully distributed system might need a trillion wires to connect all the processors with each other. But suppose that all the communication happens *through* a central processor, which further selects and processes the information to be transmitted to each of the other processors. Then, all that is needed is wires from each processor to the central one and back (figuratively called a “hub-and-spoke” architecture), which means about two million wires, enabling a reduction by several orders of magnitude. Still, the computational power of the system need not be restricted by the central processor if it is skillfully designed to “delegate” the hard computation to all the processors and only take a coordinating role. Such architectures are currently of great interest in artificial intelligence.<sup>20</sup>

In fact, the whole dichotomy between a powerful central executive and no central executive is a bit artificial. There can be varying degrees of control that a central executive is able to exercise. While it is not possible to say much with certainty on this topic, the reality in the brain may well be that there is a relatively weak

---

<sup>18</sup> (Metzinger, 2003; Eisenreich et al., 2017)

<sup>19</sup> In human neuroimaging literature, the existence of a central executive network is more or less accepted by many authors (Koechlin and Summerfield, 2007; Sridharan et al., 2008; Botvinick and Cohen, 2014; Marek and Dosenbach, 2018). The functions attributed to the central executive may vary, and many authors indeed talk about a number of “(central) executive functions” without claiming that they are performed by a single entity, whether brain region or network, as discussed by Miyake et al. (2000); Diamond (2013). For this book, the main executive function discussed is the control of actions and thoughts, as treated in the following sections; inhibition of “impulses” and automatic behaviour such as interrupts is a fundamental instance. (See Teper and Inzlicht (2013) for a discussion on how the related “self-control” could be improved by meditation.)

<sup>20</sup> For example, “federated learning” has recently emerged as such a paradigm (Kairouz et al., 2019).

central executive that controls some things to some extent, perhaps many things to a limited extent, but it does not control everything. It may be in control a lot of the time, but not when emotional interrupts, wandering thoughts, or similar processes take control of the mind. Thus, while parallel and distributed processing is inherently without central control, it may be advantageous to introduce some limited form of central executive, and this may turn out to be the best description of what happens in the brain.<sup>21</sup>

## Control as mere percept of functionality

Yet, what is undeniable is that I clearly *feel* that I can control my body and do things such as raising my arm. A central executive is often intuitively assumed to exist based on exactly such a feeling of self, or a feeling of control. But why should we assume that there is a central executive simply because it feels like there is control? The feeling of control is just another form of perception, and as we have seen, perception may not be accurate. Perception follows certain rules outlined in Chapter 10. It is usually based on incomplete information which has to be combined with prior assumptions to arrive at a conclusion, and this conclusion or inference is what we perceive. Mistakes do happen in this process.

The perception of control in the brain seems to be based on predictions—like so many other things in the brain. Every time you engage in any action, your brain tries to predict the outcome of the action. In particular, when the brain sends detailed motor commands to the muscles, it uses an internal model to predict how the limbs should move as a result. The brain then computes an error signal, comparing the predictions with the actual outcome. In humans, small errors in such predictions are actually quite common because of constant physiological changes in your muscles due to fatigue; or it could be that you are holding something heavy in your hand, which increases the force required to lift the arm. Computing the prediction errors is useful since they enable the brain to learn or adapt its motor commands to such changing circumstances.<sup>22</sup>

Now, if the prediction error is small (the actual outcome of the action is not very different from the prediction), you feel that you generated the action, and you are in control, according to current thinking in neuroscience.<sup>23</sup> This is the computational mechanism underlying the perception of whether you are in control. In contrast, if the errors are very large, the feeling of control is disturbed, and various pathological symptoms may arise. You may even feel the arm is being controlled by somebody else (by “them”, or by “spirits”), as typical of some schizophrenic patients.<sup>24</sup>

Based on his extensive psychological experiments, Daniel Wegner<sup>25</sup> proposed a related theory: the perception of control is simply based on one part of your brain observing a correlation between two things, which are the formation of an intention to act (intention being used here in the ordinary sense of the word, not in the AI

<sup>21</sup>Further theoretical neuroscience arguments on this question can be found in (Botvinick and Cohen, 2014; Rueda et al., 2004; Baumeister et al., 2007).

<sup>22</sup>(Kording et al., 2007). This is an example of the principle of feedback for successful control and action. Feedback is a general principle in action selection, which is important if there is uncertainty in the world. Just finding the best path to the goal is not sufficient if the environment is uncertain and may change. For example, if you calculate the best possible path to a restaurant, that is usually fine, and you can just walk there. But unexpected things might happen: A road might be blocked by a delivery truck; there might be construction work. This is another limitation of purely planning-based action selection, and something more similar to habit-based action selection is better in a changing, uncertain environment.

<sup>23</sup>(Haggard and Chambon, 2012; Wolpert et al., 1998; Wen and Haggard, 2018; Choudhury and Blakemore, 2006)

<sup>24</sup>(Spence et al., 1997; Frith, 2012)

<sup>25</sup>(Wegner, 2002, 2003); see also (Hommel, 2013; Pockett et al., 2009)

sense as usually in this book) and the action actually taking place. If the action happens soon enough after the formation of the intention and the action happens as you intended (and you cannot explain the action in any other simple way), the brain concludes that “you” actually performed the action out of your own “free will”. A strong correlation between intentions and outcomes is not very different from small prediction errors, and thus this psychological theory is very much in line with the neuroscience results cited above. Interestingly, just like visual neuroscientists who construct optical illusions, Wegner then devised clever experiments where the perceptual system makes the wrong conclusion about control, thus showing that the feeling of control can be fooled like any other perception.

Any of these computations are actually quite simple and could be easily implemented in a robot. A robot can assess whether it is able to control its arm by comparing the results of its motor commands and the actual outcome. Suppose some kind of central processor sends a command to the joints in the arm that the arm should be lifted by 10 cm. A couple of seconds later, the input from the camera (or input from specialized sensors in the joint) tells the central processor that the arm was, indeed, lifted by 10 cm. The central processor then concludes that it is in control of the arm.

This logic demystifies the concept of control, which is no longer anything deep or philosophical. The perception of control by the robot above is due to computations of a rather practical nature. In fact, any agent should have a model of what parts of the world it can control (e.g. its limbs) and which parts are outside of its control (e.g. mountains). This is in contrast to my everyday perception that it is *I*, or myself, that is in control, which is the result of a very complex inference process, and possibly exaggerated, misleading, or even false and illusory. Our everyday perception of control *by ourselves* is, therefore, no proof for the existence of a central executive, or “self”, that controls actions. One could say that our perception only indicates that *there is* control in the simple sense of the limbs moving as expected, but it does not necessarily mean that there is any particular entity that is in control. In other words, our feeling of control simply means that certain systems are working in a predictable way, correctly and in harmony with each other; in particular this is about the decision-making system, the motor system, and the actual limbs (or “actuators” as they are called in robotics).

### *Free will*

Free will is a celebrated and highly controversial concept in Western philosophy—the idea that you decide your actions “yourself”, that is, your actions are not merely a function of external circumstances, such as your past or other agents. Free will is very closely related to control and feeling of control: most of neuroscience uses the terms almost interchangeably. There are some nuances, though: talking about free will emphasizes your capacity to decide what you will try to do, while talking about control emphasizes your ability to actually do it, i.e. change the state of the world. A very clear difference is, moreover, that free will is almost always considered a conscious phenomenon, while control need not be, as we saw above. Even a completely unconscious robot would benefit from knowing which events are due to its own actions and which parts of the world it can control.

Philosophers have been debating about free will for hundreds of years. Democritus claimed already around 400 BCE that everything, including humans, consists of atoms, and follows strictly deterministic causal laws, thus excluding any free will. A bit earlier in India, the Buddha had debates against philosophers who held similar, strictly deterministic views.<sup>26</sup>

---

<sup>26</sup>Early Buddhist philosophy is actually often interpreted as deterministic, based on the Buddha’s emphasis on causal chains (e.g. *Samyutta Nikaya* 12.12; see also page 162 below), but he also admitted the existence of “an element or principle of initiating an action”

In a famous series of experiments, Benjamin Libet recorded an EEG response that is known to precede any action decision. The results showed that conscious experience of the decision started up to half a second *after* the beginning of the EEG response. From this, it is tempting to conclude that consciousness cannot cause the action decision, and hence there is no free will. The EEG presumably measured some unconscious processes which started the decision-making process long before any involvement by conscious processes. Libet's own interpretation, though, was that consciousness could still participate in the action decision by having the possibility of "vetoing" any decision that the unconscious circuits were trying to implement. This would imply some weak version of free will, but his interpretation is controversial.<sup>27</sup>

Some would argue that denying free will may be dangerous: people have to believe in free will in order for our moral systems to work. If people don't believe in free will, they might not feel they have any moral obligations and might behave just as they please. Our justice system in particular is based on the idea of free will: if it can be proven in court that a murderer acted without free will, say, because of a brain tumour, that will usually lead to a reduced sentence. This is, of course, not saying that there *is* free will, just that it may be *useful* to think that there is one.<sup>28</sup>

One of the most influential psychologists in the 20th century, B.F. Skinner, had a more computational viewpoint. He thought human behaviour is simply determined by rewards and punishments. That is also where the "moral" behaviour comes from; no special metaphysical beliefs are necessary. Reward people for good behaviour and punish for wrong behaviour; that is all that is needed to make them follow moral rules, in Skinner's view. From the perspective of this book, I can partly agree with Skinner on the importance of learning from the environment, where learning, again, includes evolution.

## Philosophy of no-self and no-doer

Let's go back to the "no-self" quote by the Buddha which we saw earlier (page 125). In Buddhist philosophy, it is the historical basis of a celebrated doctrine claiming that there is no such thing as the self. This is clearly a much more general idea than the mere claim that there is no entity which is in control, but in this chapter, we focus on the aspect of control and free will. We can now recapitulate the ideas in this chapter in view of justifying some claims regarding the existence of self.

First, if the brain, and thus the mind, is composed of many different processors all working simultaneously and to a large extent independently of each other, how could we speak of a self? If we admit there is no central executive, that is one form of "no-self": there is no particular part of the mind that actually is in control and could be called self in that sense. (It is not clear if it is neuroscientifically correct to deny the central executive, but let's admit it for the sake of argument here.) The *conscious* part of the mind does not control actions according to neuroscientists such as Libet and Wegner, thus contradicting our everyday perception that we decide actions on the conscious level. Decisions seem to be actually taken by various unconscious neural networks, and it may be difficult to point out any single entity making the decision.

Some parts of Indian philosophy actually formulate a more specific doctrine of "no-doer", which means that there is nobody that "does" anything in terms of taking the actions—or that at least, it is not "you" that does anything. Instead of "you" making conscious decisions and being in control, your body and mind are

(*Anguttara Nikaya* 6.38), which sounds a bit like free will. See also Federman (2010).

<sup>27</sup> (Libet et al., 1983)

<sup>28</sup> (Vohs and Schooler, 2008; Baumeister et al., 2009; Roskies, 2006)

constantly on some kind of autopilot, and your consciousness is merely observing it all.<sup>29</sup> The points above give some credence to such a variant of no-self philosophy.

But we may go further. If it is not your conscious self that decides, is it even necessarily your neural networks? In our framework, we could say that control is ultimately exercised and decisions are ultimately taken by the *input data* that the agent learns from. Our computational models have assumed that our actions are determined by past input data, together with the design of our learning and inference machines—even though the mapping from input data to action can be extremely complex and impenetrable.<sup>30</sup> From this viewpoint, nobody at all is in control, and there is no free will—while it is possible to say that *there is* some control in the specific sense of sufficient predictability. Furthermore, the existence of a central executive in the brain becomes irrelevant: if it exists, it is still only a vehicle for the evolution and the input data to steer our thinking and behaviour. Even a brain with a strong central executive could be seen as having no self: even though a central executive was seen as the hallmark of a self earlier in this chapter, if we now see all its actions as simply following from learning based on input data, it may not actually qualify as a self.<sup>31</sup>

Ajahn Brahm, a famous meditation teacher, once said that when he sits down to meditate, he always remembers the instructions of his own meditation teacher in his head; thus, it is not really Ajahn Brahm who meditates, it is his teacher—or, if I may, it is the input data he received from his teacher.<sup>32</sup>

---

<sup>29</sup>See e.g. the classic Theravadan Buddhist meditation manual *Visuddhimagga* (Chapter XIX,20); or the Advaita Vedanta teacher Nisargadatta (1982). The Buddha himself may not have formulated no-self in exactly this way; in fact, he seems to argue against it in *Anguttara Nikaya* 6.38. See Harvey (2009) for a detailed discussion.

<sup>30</sup>This is not the same as the Skinnerian viewpoint because most of the time, nobody is explicitly and purposefully feeding all the data into your brain to “train” you; most of the data is just passive observation, often with no rewards involved.

<sup>31</sup>It is obviously important to consider the exact definition of free will. In one radical viewpoint, freedom of will is a matter of not being physically or psychologically forced or compelled to do what one does. This viewpoint is called “compatibilism” in philosophy since it implies that free will is compatible with determinism (Strawson and Watson, 1998). Consider a basic robot. If it decides to raise its arm, is there any physical constraint that would prevent it from doing so? Are the computations only due to its own sensory input (in its cameras or else) and computations made in the processors inside the robot? Here, the idea of free will is formulated in terms of what causes the agent’s actions: Is it solely information-processing inside the agent—in a modern formulation—or is something outside it influencing the decision? If so, even such a robot could be said to have free will. Humans would certainly have free will. It may not be the conscious self that decides actions, but the neural networks in the brain. Still, as long as the neural networks are inside the human skull, it is, the human that decides and controls its actions. Yet, many find such a definition of free will questionable. These include all the schools in the philosophy of free will other than the compatibilists. What I described in the main text is rather similar to the “pessimist” school. Namely, an obvious counterargument to the compatibilist definition is that it depends on the time scale used: we should look back in time, trying to find the original reasons for your actions. As I argue in the main text, fundamentally, the robot’s or human’s actions are just a result of its programming/evolution and, especially, the input data, so there goes free will.

<sup>32</sup>This example also points out how a lot of the input data comes from social interaction, which is outside of the scope of this book.

## Chapter 12

# Consciousness as the ultimate illusion

Why do we have conscious experiences? This is one of the deepest unanswered questions in modern science. It is not even quite clear what the whole question means and how it could be formulated in a rigorous, scientific manner. One thing that is clear, however, is that consciousness is somehow related to suffering. Some would even claim that in a strict sense, there can be no suffering without consciousness.

In this chapter, I try to shed some light on the nature and possible functions of consciousness. I consider two different aspects of consciousness: it can be seen as performing some particular forms of information-processing, or it can be seen as a subjective experience. I provide a critical review of the main theories concerning these two aspects. In particular, I explain how consciousness is related to mental simulation and the self, and as we have seen, those play an important role in suffering. This leads to some old, but still radical, philosophical ideas about the nature of our knowledge of the world and how it relates to our consciousness. Ultimately, I argue how changing your attitude to consciousness may actually have a strong influence on your suffering, a theme that will be further elaborated in later chapters.

### Information processing vs. subjective experience

The main problem we immediately encounter in research on consciousness is the difficulty in defining the terms involved. “Consciousness” has different meanings to different people and in different contexts.<sup>1</sup> For our purposes, we can divide the concept of consciousness into two aspects. First, there is the information processing performed by human consciousness. This is something we might understand based on AI, since information processing can usually be programmed in computers. One approach is to ask what the computational function, or utility, of consciousness might be in humans; these are relatively well-defined scientific concepts and questions. This approach is fine as long as we are content to consider consciousness as another form of information-processing, or computation. The second, more difficult aspect of consciousness is the experience. That is, the conscious “feeling” which is specific to myself, i.e., subjective. Its existence is so obvious that it is rather neglected by most people.

When you look at the text in this book, several quite amazing things are happening; they can be roughly divided to information processing and experience. Those related to information-processing have been discussed earlier in this book. Light enters your eye, generates electrical signals on the retina, the signals travel into your

---

<sup>1</sup> (Van Gulick, 2021)



brain, and some incredibly intricate information processing takes place, allowing you to recognize the letters and even transform the letters into words. However, all that is simply information processing, and it can soon be programmed in a computer—in some more rudimentary form, it is possible even now.

But in addition to such information-processing, there is something else: you have a conscious, subjective *experience* of the book, the letters, and the words. Somehow, almost magically, the book appears in some kind of a virtual reality created by your brain. We tend to think that this is normal since the book is there, and we simply “see the book”. But in fact, the conscious experience is not somehow *in* the book, and it does not somehow automatically come *out of* the book. The experience, the awareness of the book is created by some further mechanisms which we simply don’t understand yet. This experiential aspect is called “phenomenal” consciousness. Philosophers use the word “qualia” in this context: the conscious “quality” of the book being seen, “what it is like” when the book is consciously experienced. Or, as more poetic narratives would have it, it is the “redness of a rose”. It is not information processing but something more mysterious.

It is this phenomenon of subjective experience, or qualia, which is the main topic of this chapter. It is also the main meaning in which I use the word “consciousness” in this chapter; “awareness” is used in exactly the same meaning, and so is “conscious experience”.

## The computational function of human consciousness

Now, what is the connection between these two phenomena: information-processing and consciousness? Conscious experience is certainly not just one form of information-processing, but the connection is extremely difficult to understand. In fact, consciousness must have *some* connection with information-processing: the qualia of the rose must be based on processing of incoming sensory input, even if most of that sensory processing seems to be unconscious. Let us assume, in the following, that part of the information-processing in the human brain is conscious, in some sense to be elucidated. This is such a typical assumption that it is often not even made explicit.

Let us then try to understand what we can say about the function, or utility, of such conscious information-processing. Taking a more neuroscientific approach to the question, one can first ask: What are the evolutionary and computational reasons why certain animals, such as humans, have consciousness? We assume here that consciousness is a faculty that is a product of evolution—but strongly influenced by culture, of course. It is quite difficult here to ignore the experiential part of consciousness and consider information-processing only. If we say that an animal, or an AI, is conscious, it seems to almost necessarily mean a conscious experience: we wouldn’t even know what it means to say that an animal is conscious if it does not have conscious experience. So, in a sense, the question is necessarily about the *computational function* of human conscious *experience*, and whether it can be explained by evolutionary arguments.<sup>2</sup> I will next review a number of proposals.

Investigating wandering thoughts actually lead us close to consciousness, because “thinking” is often considered the hallmark of consciousness. More precisely, the fact that we can reconstruct a vivid image in our minds about past or future events, while ignoring the present sensory input, is a remarkable property that seems to be closely related to consciousness. Some investigators actually propose that one of the main functions of consciousness is such *simulation*, which is also called *virtual reality*. That is, consciousness allows

---

<sup>2</sup>For discussions on the consciousness in general and the computational function in particular, see (Baars, 1997; Chalmers, 1996; Seth, 2009; Van Gulick, 2021).

us to consider different scenarios of what might happen in the future, and what would be the right things to do in those circumstances. Planning crucially needs a capacity of simulating results of future actions, and indeed in the case of wandering thoughts, we already talked about simulation. Such simulation obviously would be useful for survival and reproduction, and thus favoured by evolution. A special case of such simulation is dreaming, which does create a virtual reality that is particularly far removed from current reality. Dreaming often includes simulation of threatening situations, i.e. situations in which it is important to know what to do to avoid harmful consequences.<sup>3</sup>

However, I think we should not too easily conflate consciousness with thinking or simulation. What we see is a correlation between a certain brain function (namely, simulation) and consciousness, but it is difficult to say whether consciousness is really essential for such a function. As we saw above, modern AI uses planning, and even systems similar to wandering thoughts, simulating events from the past and simulating episodes that might happen in the future. Yet, nobody seems to claim that replay or planning would make a computer conscious. Such a claim seems absurd to most experts because claiming that a computer is conscious is usually interpreted as having phenomenal conscious experiences; but they are very unlikely to be produced by such simple computations as replay and planning.

Another possible function of consciousness is *choosing actions*. We typically have the feeling that we consciously decide what we are going to do, an experience of free will. You may think that you decided to read this book; perhaps you decided to read this particular sentence. But did you actually decide how you move your eyes from one word to another? What do we actually decide on a conscious level? As we saw in Chapter 11, consciousness may not have *any* role in the control of actions; the feeling of free will and control may be deceiving. It may very well be that actions are entirely decided by unconscious processes. After all, that seems to be the case with many animals (if we assume most of them don't have consciousness), as well as any robots and AI that exist at the moment.<sup>4</sup>

Yet another proposal is that consciousness could be useful for *social interaction and communication*.<sup>5</sup> The contents of consciousness can usually be communicated; in fact, in psychological experiments, one operational definition of the consciousness of perception is that you can report the perception verbally to the experimenter. The utility of the conscious perception, in particular, would be that this perception can be transformed into a verbal form, and communicated to others. Again, a problem is that it is perfectly possible to build AI and robots which communicate with each other without anything we would call consciousness, at least in the experiential sense.

A particularly relevant proposal for this book is that consciousness facilitates *communication between different brain areas*.<sup>6</sup> While unconscious processing has a huge capacity for information-processing, it suffers from the problem that the processing is divided into different brain areas whose capacity for communication is limited—as typical of parallel distributed processing. The idea here is that consciousness is the opposite: it has very limited processing capacity, but its contents are broadcast all over the brain. Consciousness can thus be considered a “global workspace”. It could be compared to a notice board where you can put short notes (limited capacity), which will be seen by everybody in the office (global broadcasting). It is also a bit like a cen-

---

<sup>3</sup>(Revonsuo, 2000; Hesslow, 2002)

<sup>4</sup>For a viewpoint in which consciousness is the “reason” for most actions, see Cleeremans and Tallon-Baudry (2021).

<sup>5</sup>(Hommel, 2013; Frith, 2002, 2010)

<sup>6</sup>(Baars, 2002, 1997; Dehaene and Naccache, 2001). Partly related to this are the “higher-order” theories in which consciousness depends on higher-order mental representations that represent oneself as being in particular mental states (Lau and Rosenthal, 2011), that is, a form of metacognition.

tral executive in the society of mind discussed in Chapter 11—one which is not particularly smart but whose thundering voice is easily heard even at a great distance. This links clearly to the proposal in previous chapters, where we considered a system where pain, or errors such as reward loss, are broadcast to the whole system. An intriguing possibility is that this could be why pain, whether physical or mental, must be conscious. Perhaps pain is so acutely conscious precisely because the broadcasting system it uses is inherently related to the global workspace of consciousness. Yet, it is not clear to me why this would require conscious experience, since distributed information processing is increasingly performed in computers as well.

I have just described several proposals which each consider highly relevant information-processing principles. For example, inside an AI, the communication between different processors or processes needs to be solved, and mechanisms related to the global workspace theory can be very useful. Yet, in each of those cases, we have to ask whether would we say an AI with such capacities is conscious. Would it necessarily have subjective experience, if that is what we mean by “conscious”? We could go through *all* the computational functions of the preceding chapters and ask whether consciousness necessarily has any role in any of them. To the extent that all of these are simply computations that can be implemented in an AI, they may actually not *need* consciousness. (I’m here assuming that any computer we have at the moment is not phenomenally conscious, which is relatively uncontroversial.) Therefore, I think there is currently little reason to believe that consciousness would be necessary for some particular kinds of computation, which would be impossible without consciousness.

### *Consciousness as a specific hardware implementation?*

However, another viewpoint is possible: there may be some forms of information processing which are *correlated* with consciousness. It could be that some of the computational routines in the brain are always implemented using some special circuits or processes that give rise to consciousness. Such computations would then always give rise to consciousness, even if in theory, it would be possible to implement them in non-conscious circuits. If we program that same kind of computation in an AI, we might then say that we have programmed the AI to perform “conscious” information-processing. However, it may be best to use scare quotes here: the AI may be imitating processing that is conscious in the brain, but it might not have conscious experience, so whether we should call such computations “conscious” is questionable.<sup>7</sup>

Therefore, any argument—such as I have just made—saying that a computational function cannot be the actual function of consciousness because it can easily be programmed in an AI, may be missing the point. While it may not be completely necessary to have consciousness for, say, simulation, it could still be that in biological organisms, shaped by evolution, consciousness is somehow an important part of the computational implementation of simulation, or any of the other functions above. The fact that something is easy to program in a computer, which is based on completely different kind of hardware, does not mean that it might not be very difficult to implement in the brain without the help of some, hitherto unexplained, conscious mechanisms. Thus, consciousness might be a particular “hardware” implementation of certain computations which

---

<sup>7</sup>The well-known distinction between “access” and “phenomenal” consciousness (Block, 1995; Kouider et al., 2010) is related to this point. In this book, when I talk about consciousness, I mean phenomenal consciousness, i.e. the experiential kind of consciousness, unless otherwise mentioned (or in quotation marks). Access consciousness is, in my view, an operational definition of consciousness, used in experimental neuroscience: If you ask a person whether she is conscious about X, and they reply yes, then the person is conscious of X in the sense of having access to the experience or perception of X. I find this definition of access consciousness not very relevant for the present discussion.

are otherwise difficult to perform in the brain—whether simulation, global workspace, or else.

Yet, this is all mere speculation. We cannot exclude the completely diametrically opposed possibility, which is that consciousness is not actually part of the information processing at all. Perhaps it does not affect the computations, or the contents of the mind, in any way; it simply *reflects* the results of the information processing. It might not even have any evolutionary utility.<sup>8</sup> Many thinkers over the centuries have proposed that consciousness in humans is only the tip of the iceberg, and most mental activities—which I call information processing—happen without consciousness.<sup>9</sup> But here, we find an even more startling possibility: perhaps consciousness is not even the tip of the iceberg but, to push the metaphor further, a bird that flies over the iceberg, only watching it from a distance. We will see even more startling possibilities later in this chapter.

## The origin of conscious experience

Next, let us consider the problem of the existence of conscious experience itself. Most scientists would agree on the fundamental importance of understanding the physical, chemical, and biological processes that enable conscious experience. While it is one of the deepest questions in science, I'm afraid there is little we can say about it with any certainty. It is not even clear if the whole question can be approached scientifically. This is because it is difficult to make any rigorous experiments on experience because it is subjective. Only I observe my conscious experience; you, or any neuroscientist, cannot really know what I experience. So, how could a neuroscientist conduct experiments on people's experiences? Measuring brain activity, or looking at people's behaviour are not measuring experience. Brain activity and behaviour are related to and correlated with experience, but not the same thing. The closest you can get is asking people what they experience. However, they might not be able to express it verbally with sufficient accuracy or detail. In fact, if participants in experiments answer such questions, they are ultimately engaged in behaviour (in the form of speaking), and, in a sense, the neuroscientist is actually only measuring their behaviour (speech, in this case).

With good reason, the problem of understanding how and why the brain creates conscious experience—including whether it is actually the brain that does that—is called the *hard problem* of consciousness research.<sup>10</sup> However, let us not despair: Even if any solution may not be available, some interesting things can be said about the problem.

To begin with, some neuroscientists think there is something special in humans that enables consciousness, and perhaps in some other mammals such as great apes as well. What it would be, nobody really knows. The main theories are based on observing what kind of structures human brains have, and what simpler animals like cats and dogs do not have. Because the brains of cats and dogs are in many ways very similar to the

---

<sup>8</sup>I'm not committing myself here to any particular philosophical stance, but just exploring different possibilities. Claiming that consciousness has no evolutionary function could be seen as a special case or implication of the stance called Epiphenomenalism (Walter, 2009), according to which consciousness has no causal effect on physical events. What I have written here is probably also compatible with the sophisticated alternative given by Chalmers (1996), who also addresses the obvious question of why there might be consciousness if it is not evolutionarily advantageous (his Section 3.6). Chalmers's arguments rely heavily on considering what a zombie without any consciousness would be like compared to humans; I think his zombies are comparable to AI, which I assume is not conscious. Chalmers (1996) seems to agree on the difficulty of understanding consciousness: "[W]hen it comes to consciousness, it seems that all the alternatives [of philosophical stances] are bad. If someone comes away with the feeling that consciousness is simply an utter mystery, then that is not completely unreasonable."

<sup>9</sup>See footnote 5 in Chapter 7 for some historical remarks.

<sup>10</sup>(Chalmers, 1995, 1996)

human brain, the relatively small differences might be related to consciousness.

One difference between the brains of humans and “lower” animals seems to be the existence of a special class of neurons, called von Economo neurons. They have particularly many long-range connections to other neurons. Since long-range connections might be related to something like a global workspace, von Economo neurons have often been considered as a potential candidate for a mechanism generating consciousness.<sup>11</sup> In fact, it has also been suggested that the brain basis of consciousness might be related to feedback between brain areas, as opposed to any special kind of processing inside each single area.<sup>12</sup> It could be that the long connections of von Economo neurons make such feedback stronger, sufficiently complex, or otherwise more conducive to consciousness. Interestingly, apes have von Economo neurons as well, and so do elephants, dolphins, and even some monkeys, so based on this criterion, those animals at least should be conscious.<sup>13</sup>

Alternatively, we can use AI for studying the hard problem of consciousness. We can perform thought experiments, based on the same kind of comparison as was just done with other animals. Assuming that AI is not conscious, part of the hard problem of consciousness is then to explain what creates this fundamental difference between humans and AI.

### *How can we know something is conscious*

However, there is a problem with the argumentation above: It is based on finding animal species, such as cats and dogs, which are reasonably intelligent but have no conscious experience. Or, if we consider AI, it is assuming that AI is not conscious. But how can we even know if an animal species or AI is conscious or not?

Some would claim that we cannot even know if other people are conscious. We do tend to assume that every human we meet is conscious, but this is just a guess, really, without much logical basis. We are actually generalizing based on ourselves: the only human I know for sure to be conscious is myself. Others just move around and say things, but they could be some kind of robots for all I know; perhaps I am the only person conscious in the world. If I assume all other humans are conscious as well, I can only hope I’m not overgeneralizing! This is, somewhat cheekily, called the zombie problem: it could very well be that some of the people you meet are “zombies”, that is, creatures that look like humans and behave like humans, but do not possess any kind of consciousness.

Leaving such wild speculation aside, we do have a real scientific problem here. In neuroscience, it has been found extremely difficult to determine which animal species are conscious and which are not.<sup>14</sup> Even considering humans, it is not easy to tell whether people in coma are conscious. For coma patients who are incapable of saying anything or any motor responses whatsoever, measuring brain activity provides the last resort for assessing their consciousness. Surprisingly, it has been found that patients who were thought to be

<sup>11</sup> (Critchley and Seth, 2012; Butti et al., 2013)

<sup>12</sup> (Lamme and Roelfsema, 2000; Crick and Koch, 2003)

<sup>13</sup> It is not very clear which animals have von Economo neurons and which don’t; Jacob et al. (2021) have recently claimed to find them even in raccoons.

<sup>14</sup> (Seth et al., 2005). Neuroscientists have developed an interesting test called Mirror Self-Recognition (Toda and Platt, 2015). The idea is that a mirror is introduced to the animal. After the animal has had some experience with the mirror, some red dye is applied to its face to create a small but visible spot. Many animals instinctively try to touch the red spot. But does the animal touch the real spot on its face, or its image in the mirror? If it touches the real spot, it is concluded that the animal had some kind of consciousness of itself, or at least a body image similar to what we have. Chimpanzees, for example, pass the test. However, this is of course a very indirect measure of only one aspect of consciousness, in particular self-awareness considered later in the text. (For moral implications of our ignorance of whether animals can suffer on a conscious level, see Birch (2017).)

in a completely unconscious, vegetative state are sometimes perfectly conscious, in the sense of being able to respond to questions like healthy humans would.<sup>15</sup> Such people can sometimes learn to communicate with the external world by special devices which transform brain activity to text. So, it is actually true that we cannot always tell if even other humans are conscious!

Likewise, how could we then judge whether an AI is conscious or not? What if current AI is conscious, or will become conscious in the near future? You can find people arguing strongly for the possibility of conscious AI. Some say it is simply a question of complexity: when AI becomes complex enough, it will become conscious; the only reason why present computers are not conscious is that they are too simple in terms of their computation, in particular lacking sufficient interaction and information interchange between different processing units. Others think an AI must have a body, i.e. it must be a robot, in order to be conscious, and consciousness is somehow created in the interaction with the world.<sup>16</sup>

Fundamentally, the question of determining consciousness seems to be unsolvable because of its subjective nature: I can only know something about my own consciousness. We cannot know for sure if any animal or AI is conscious or not. Consciousness—at least regarding its experiential quality—remains a huge mystery.<sup>17</sup>

## Why is simulated suffering conscious?

Let's get back to the question of suffering. Consciousness is in some sense crucial to suffering: if we were not conscious of our suffering, if we didn't have the conscious experience of suffering, it would not be the same kind of suffering at all. Suppose you have a headache but you start watching a really fascinating movie; you may cease to notice the pain at all. Somewhere in your brain there is probably some kind of activity which would usually lead to the experience of pain, but your attention is in the movie, so you completely ignore the pain. That is because when you are not paying attention to something, you cannot be conscious of it either.<sup>18</sup> So, in some sense, the whole problem of suffering revolves around the question of consciousness. If we consider a simple animal or an AI and agree it is not conscious, is it actually meaningful to say it suffers—as I may have done in this book?<sup>19</sup>

The other day I was watching a fictional TV series in which a tiger attacked a woman. I felt scared. Was there

---

<sup>15</sup>(Monti et al., 2010; Bruno et al., 2011)

<sup>16</sup>For reviews, see (Reggia, 2013; McDermott, 2007; Chella and Manzotti, 2007); on complexity, (Tononi and Edelman, 1998), and embodiment, (Ziemke, 2007)

<sup>17</sup>However, I don't mean to be completely pessimistic about the possibility of doing scientific research on consciousness. If verbal reports (or similar information) of conscious content are combined with brain imaging in a sufficiently large number of human subjects—possibly specifically trained to perform introspection—progress can be made. The specific methodology needed is discussed by Lutz and Thompson (2003); Gallagher (2003).

<sup>18</sup>The connection between attention and consciousness is complex, but it is usually assumed that we can only be conscious of something we attend to. De Brigard and Prinz (2010) review evidence for and against this assumption.

<sup>19</sup>This problem could be contrasted with the problem of whether a computer can *see*. Suppose a robot moves around in its environment, avoiding obstacles and performing some task thanks to input from its camera. Now, how would you answer the question of whether the robot is able to “see”? Most people, including scientists working on such computers, would casually say that the computer sees, for example, it sees the obstacles. If pressed hard on what that means, they would probably admit that the computer “does not *really* see”, presumably because there is no consciousness involved. What is very interesting is that this ambiguity is not usually considered a problem: it is rare that any serious debates are conducted on whether such a robot actually “sees” or not. When we talk about suffering in an AI, the situation is, in principle, quite similar. However, much more heated debates can be expected on the question of whether the AI actually suffers. This lack of clarity on whether an AI can suffer seems to be much more difficult to accept than in the case of seeing.

any point in being scared? I was in my own home, just watching an electronic device produce some patterns of light on its screen. There was no real tiger near-by, no real risk of being eaten. Even if I had been in the middle of the action, it would have been on a film set. The tiger was tame; or perhaps it was just a computer animation, and there was no real tiger at all. In any case, even if I had been at the studio instead of home, I would not have been in any kind of physical danger. What is even more interesting is that after having watched that on TV, my brain started replaying the events. Several times during that evening, I saw the tiger in my wandering thoughts. Every time some element of fear crept into my mind. I thought: How stupid can my brain be? Why do I feel fear, although there is no real tiger near-by, there was never any real danger of anybody being eaten by a tiger, and finally, I haven't even seen the image of a tiger for hours, it's just repeating in my head.

This is yet another amazing thing about conscious simulation: It reproduces the same valences, that is, the positive and negative feeling tones, and the same experience of suffering, as the real thing. When I think about something unpleasant, it hurts. Maybe not quite as much as the real thing, but still it hurts. The theories explained in previous chapters actually explain, to some extent, why the brain does that. It is not stupid to replay experiences. Replay and wandering thoughts are important for learning a good model of what the world is like and what kind of actions are useful in which situations, as we saw in Chapter 9.

Yet, my current accusation of my brain being stupid is on a different level than the theories of the previous chapters. Here, I'm talking about consciousness. Why am I *consciously* afraid of the tiger, and consciously suffering during the replay? Why do I need to experience suffering while the brain is performing such simple computations that we can easily program in a computer? To put that more precisely, why do I need to experience a negative valence on a conscious level while doing the replay? Couldn't the brain just do the replay somehow quietly on an unconscious level without disturbing my conscious feelings and conscious thinking? So, I'm not just repeating the question of Chapter 9, which was: why do wandering thoughts trigger feelings of pain and pleasure. Here, I ask a more general question about consciousness and suffering: Why are such simulations, and the ensuing suffering, *conscious*?

Again, we might assume that perhaps evolution just made a simple computational shortcut. If something dangerous is perceived in the outer world, the fear system has to be activated on every level, unconscious as well as conscious. In particular, one possibility is that conscious fear is important in the information-processing because of its capability for broadcasting as in the global workspace hypothesis. Therefore, conscious fear has to be activated to properly compute things.

There is another possibility. Above, we saw that consciousness might not be necessary for any computations, which would invalidate the argument just given. Now, assuming conscious fear has no computational utility, it could still be the case that it would be too much trouble to somehow switch consciousness off when doing replay. It would be nice indeed if, when something dangerous comes up in a wandering thought, the fear system would be activated only partly, not on the conscious level, perhaps only in some distant corner of the unconscious processing systems. This would be nice, but would evolution have any reason to do us such a favour? We should recall again that evolution does not care at all about whether we feel good or bad. It tries to optimize computation in order to maximize the spreading of the genes, and this has to be done with limited computational resources. Allowing us to switch off conscious suffering when engaging in replay would presumably be pointless from the viewpoint of optimizing computation. So, evolution just makes us suffer from replay since that is optimal use of finite computational resources. Such optimization of computation may actually increase our chances of survival a bit, and give us a longer life. Full of suffering, though.

## Self vs. consciousness

So far, I have been considering consciousness on the sensory level, as in “consciousness of the text you are seeing”. Another very different thing that we can be conscious of is our own self. It can even be argued that if there is any consciousness at all, there must necessarily be self-consciousness, or self-awareness, that is, conscious experience related to oneself. It can be seen as a particularly automatic and primitive form of consciousness. So, we find yet another meaning for the term “self”—in addition to those in Chapters 6 and 11— defined as precisely this self-awareness. This corresponds very well with our intuition, where it is my conscious feeling of being “me” that defines what “I” am, or what my “self” is.<sup>20</sup>

This aspect of self-consciousness is very different from the way “self” was treated in previous chapters. The aspects of self treated in earlier chapters do not necessarily have anything to do with consciousness: all the operations described earlier are just computations. In particular, an AI does not need to be conscious to infer that it can control certain things and not others, or to develop behavioural mechanisms that ensure its survival, while even a simple AI system can and should have methods for evaluating the performance of “itself”.

If self is defined in this sense of self-awareness, it might in fact be difficult to defend any form of “no-self” philosophy, of which we have seen one version in Chapter 11. Descartes famously was absolutely certain that he could say “I am” because he “thinks”:<sup>21</sup>

[A]fter having reflected well and carefully examined all things, we must come to the definite conclusion that this proposition: I am, I exist, is necessarily true each time that I pronounce it, or that I mentally conceive it.

Yet, Descartes was quite wary of saying *what* he actually is:

I must be careful to see that I do not imprudently take some other object in place of myself, and thus that I do not go astray in respect of this knowledge that I hold to be the most certain and most evident of all that I have formerly learned.

The complexities of no-self philosophy largely come from the tension between these two viewpoints: It is intuitively clear *that* I am, but it is not clear *what* I am. (There can hardly be any difference between the “I” and the “self”, they are just two words for the same thing.<sup>22</sup>)

On the other hand, some would say that such self-awareness can be seen as a mental construction, even an illusion, just like control and free will. Our self-awareness could be based on a collection of the awarenesses of various sensory perceptions, with no special core that could be called “me”, or awareness of myself. Hume expresses this potently in a famous quote which is not unlike anything Buddhist philosophers might have said:<sup>23</sup>

For my part, when I enter most intimately into what I call myself, I always stumble on some particular perception or other, of heat or cold, light or shade, love or hatred, pain or pleasure. I never can catch myself at any time without a perception, and never can observe any thing but the perception. When my perceptions are removed for any time, as by sound sleep; so long am I insensible of myself, and may truly be said not to exist.

<sup>20</sup> (Dennett, 1992; Gallagher and Shear, 1999; Gallagher, 2000; Smith, 2017).

<sup>21</sup> *Meditations on First Philosophy*. That was, in effect, the *only* thing he could say with any certainty.

<sup>22</sup> (Smith, 2017)

<sup>23</sup> (Hume, 1739), Section 1.4.6



This suggests a no-self philosophy where self-awareness is nothing but a complex of various instances of sensory awareness, mistakenly leading to an illusory perception of a separate entity called “self”.<sup>24</sup> In the absence of any perceptions, ultimately I may be said not to exist. Such no-self philosophy could be called ontological: it claims that the self does not exist at all. It does not merely say that self is not what it looks like, or that it is missing something, or that it is not too important; instead, it claims that self does not exist, period. While Hume may not have meant to go quite that far, many Buddhist philosophers do.<sup>25</sup>

## Nothing is real?

Saying that the self is a mental construction, possibly an illusion, sounds quite radical. Well, how about going a bit further, and denying that *anything* really exists? While it is undeniable that there is some kind of experience of the world outside of myself, it is equally undeniable that this experience is not the same thing as the world outside. The conscious experience is—according to a conventional neuroscientific view—the product of complex information-processing of incoming signals. Actually, most of conscious experience has little to do with the world that surrounds us here and now, since conscious contents are often a product of planning, replay, and other kinds of thinking and imagination. The interesting thing is how people are misled into believing that this experience, this virtual reality, this simulation, replay, or planning, is actually the reality.

It should be easy to admit that when we plan the future, the planned events are just imagined, and not real. But the “unreality” of consciousness goes deeper than that: In fact, *everything* in our consciousness is a simulation, a virtual reality, constructed by our mind. This also includes your consciousness of everything you see, hear, feel, taste, and smell at this very moment. Any perceptual experience, as well as any thought, is simulation, or computation, and not the same as reality.<sup>26</sup>

This is just a rephrasing of well-known neuroscientific facts. As we have discussed several times by now, when you look at this text, your brain is doing complex computations based on the incoming information. Based on the results of those computations, it creates a conscious perception, which contains an image or a feeling of the world around you, including the book or the computer screen on which you see this text. The

<sup>24</sup>In addition to the obvious sensory modalities, it is important here to consider proprioception (perception of the body position, including body ownership, Tsakiris et al. (2007); Seth and Tsakiris (2018)) and interoception (sense of the internal state of the body, in particular internal organs, Craig (2009)). Further related phenomena include meta-awareness (discussed later in Chapter 14) and autobiographical memory (Prebble et al., 2013).

<sup>25</sup>While the no-self philosophy is widely associated with Buddhism, different Buddhist schools actually approach it in very different, even contradictory ways, and several interpretations exist. In fact, the philosophy of no-self has perhaps as many facets as the very concept of “self”. We already saw the interpretation of no-self as lack of control in Chapter 11. Another approach is to see “no-self” as a suggestion not to worry about self-evaluation or self-preservation, which was the interpretation of “self” in Chapter 6; rumination may not be possible without some concept of self to which the bad things are happening. In this latter sense, it may not be so much a “truth” describing the world, but rather a useful way of thinking, as we will see in Chapters 14 and 15. The ontological interpretation we have in this quote by Hume is yet another approach, probably the most well-known in Buddhist philosophy. At the risk of greatly oversimplifying this complex issue, I would venture to say that the Theravadan school is more in line with Hume here; Theravada considers self as an illusion, as something that does not exist. In contrast, Mahayana schools, with the possible exception of Madhyamaka, emphasize the primacy of consciousness, like Descartes, and do not deny the existence of self—although they do point out that our ordinary conception of self is mistaken in various ways. See Verhaeghen (2017) for a short, readable overview emphasizing some practical implications of such a philosophy; Vago and David (2012) emphasize how mindfulness works largely through self-related mechanisms. Harvey (2009) and Williams (2008b) give book-length expositions of the philosophy.

<sup>26</sup>Note that the discussion regarding simulation in this chapter has nothing to do with the idea that we would be living in a simulation programmed by some other race, sometimes called the “simulation argument” (Bostrom, 2003).

conscious experience is created by some quite quasi-miraculous mechanism, which science has not yet been able to explain—even saying that it has to be in the brain is speculative. But the important point is that what you see is the virtual reality, or the simulation in the brain, not the real world. The distinction between the world and your conscious experience is basically inherent in the very notion of “experience”. Although I have already said this in the beginning of this chapter, this point requires a longer explanation, so let me try.

Usually, you would say that you “see” this book (let’s just assume for the sake of simplicity that you are reading this text in a book). However, according to the conventional neuroscience viewpoint, what you’re actually conscious of is the interpretation created by your brain, not the book itself. The book simply reflects some photons emitted by a lamp or the sun, these photons enter your eye, and your eye sends electrical signals to your brain. Based on these electrical signals combined with the prior information about the world, your brain creates a virtual reality, including your perception of this book. Meanwhile, based on other sensory information, and again all kinds of internal information and processing, the brain creates your perception of your surroundings, your body, and indeed, your perception of your self.

I am not denying here that the book exists. I am merely pointing out that your consciousness, your sensory awareness of this book and everything else is created by your mind, presumably by some highly complicated process in your brain. You cannot really “see this book”, you cannot be “conscious (or aware) of the book”, you are only aware of the results of some computations performed in your brain, in which the book only plays the role of being the physical source of some radiation which was input to the computations.

The metaphor of virtual reality means that consciousness is similar to wearing virtual reality goggles which feed an input to your eyes which is so realistic that it looks almost like real. In the case of seeing this book, though, it looks *exactly* like real to you because you know nothing better: You have never seen anything which would be somehow closer to reality than this virtual reality. A number of science fiction movies are based on the idea that somebody could feed fake sensory information directly to your brain, and you would have no idea the sensory input is fake. Descartes already proposed that he cannot trust his perception because an “evil demon” might be feeding an illusory external world to his mind—which is precisely why he could only be certain of his own existence. Such claims lead to an extreme form of uncertainty regarding perception.<sup>27</sup>

My point is that something like that is actually happening to you all the time, according to perfectly mainstream neuroscience. I want to emphasize that I’m not trying to make some radical philosophical point here. There are others that will tell you that the world does not *really* exist, including proponents of some Eastern philosophical systems, such as Advaita Vedanta, or Mahayana schools of Buddhism, including Zen and Yogācāra.<sup>28</sup> I’m trying to steer away from such philosophical speculations about what exists, and merely point out some of the limitations of our perceptual and cognitive systems, in a way which is, I hope, acceptable, even

---

<sup>27</sup>Descartes’s *Meditations on First Philosophy*

<sup>28</sup>For example, Williams (2008b, p. 94) describes the Yogācāra viewpoint by Vasubandhu as “Apparently external objects are constituted by consciousness and do not exist apart from it. (...) There is only a flow of perceptions.” Claiming that the world does not really exist is a form of *ontological* idealism, while claiming that we cannot possibly know for sure if the world exists is *epistemological* idealism (Guyer and Horstmann, 2018). In Mahayana Buddhist philosophy and, especially, its Western commentary, there has been a lot of debate on which form to support. For example, Lusthaus (2013) warns about misunderstanding the Yogācāra literature to constitute an ontological statement while it is actually intended to be epistemological only. (For my part, I’m not committing to any such philosophical viewpoint here.) Even in early Buddhist texts you find claims related to such idealism: “And what, bhikkhus, is the all? The eye and forms, the ear and sounds, the nose and odours, the tongue and tastes, the body and tactile objects, the mind and mental phenomena. This is called the all.” (*Samyutta Nikaya* 35.23). However, this formulation open to interpretation and may also be seen as admitting the (ontological) existence of outside objects.

if unpalatable, to most scientists working on those topics.

At the risk of repeating myself: Most neuroscientists would agree that sensory processing in the brain is producing an interpretation of the incoming input; they would further agree that the brain creates consciousness. Thus, the contents of consciousness are not a direct product of the world, let alone the same thing as the world; it is a construction, an interpretation created by the brain. Yet, we often have the intuitive feeling that the contents of consciousness are somehow identical to the contents of the outside world, which is not the case. Just studying an introductory course in neuroscience or in AI might be enough for many people to give up such an idea. Visual illusions, such as in Fig. 10.4 on page 118 are one way of demonstrating how perception is different from reality.

The Belgian artist René Magritte has a famous painting called *La Trahison des images*, or “The Treachery of Images”.<sup>29</sup> The painting consists of a picture of a pipe, with the text “Ceci n’est pas une pipe”, or “This is not a pipe”, written underneath it. The point is that the painting is just a picture, not the real pipe. While the artist’s purpose was to illustrate the deceiving nature of images, the painting illustrates the illusory nature of consciousness as well. Suppose you actually hold the pipe in your hand and look at it. What appears in your consciousness is a picture, a simulation, or a reflection of the pipe; it is not a pipe. Yet, we have the habit of thinking that the perceptual image is the real pipe, while in reality, it is only somehow indirectly related to the real pipe. Furthermore, the category of a “pipe” is just a mental construct. In this sense, perception is not the real thing; consciousness is not the reality.<sup>30</sup>

These philosophical points are not simply theoretical speculation: Our attitude to consciousness has a direct effect on suffering. Consider the example of the tiger I saw on TV: If I could somehow develop a different attitude towards the contents of my consciousness, seeing them as mere simulation, I might suffer less. This is precisely why some Buddhist schools claim that the outside world only exists in your imagination—or at least they recommend adopting such an attitude towards the world.<sup>31</sup> In the next chapters, we will consider this and many other ways of reducing suffering by changing our thinking patterns as well as using meditation.

<sup>29</sup>See e.g. <https://en.wikipedia.org/wiki/File:MagrittePipe.jpg> which cannot be reproduced here for copyright reasons.

<sup>30</sup>In this chapter, I have taken the viewpoint of physical materialism by assuming that the pipe actually exists and that our consciousness is created by the brain. If we reject one of these assumptions or both, the conclusions will of course be even more radical.

<sup>31</sup>Any such philosophical claims in Buddhism could, in fact, be seen as clever devices only intended to help with meditation and other practices (Schroeder, 2004). The Theravadan master Ajahn Chah seems to have this intention when he says “If you think things are real there is suffering and there is fear. You are afraid of the different ways things may turn out. (...) There is thinking, then fear follows immediately. It deceives you, creating a picture to mislead you. (...) As to what is actually happening, there is nothing” (Chah, 2001). See footnote 28 above for discussion on the philosophical claims concerned.

## **Part III**

# **Liberation from suffering**

**The final part will describe methods for reducing suffering, largely drawing from philosophical traditions such as Buddhism and Stoicism, while showing how they logically follow from the science of Part I and Part II**

## Chapter 13

# Overview of the causes and mechanisms

In this final part of the book, we move to the question of how to reduce suffering, or, ultimately, how to be liberated from it. Applying the scientific theories of the previous chapters, we devise various methods to that end in the following two chapters. But first, in this chapter, I will recapitulate the basic theory. I will use two flowcharts to illustrate the basic mechanisms of suffering: the first one emphasizes adverse properties of the world and general cognitive design principles, while the second flowchart focuses on the dynamics of moment-to-moment cognition. The flowcharts also make some explicit connections to the basic concepts of Buddhist philosophy. The connection of reward loss to the whole architecture of intelligence is then succinctly summarized in one single “equation”, which directly suggests ways of reducing suffering.

### Why there is (so much) suffering

Let us start by looking at the difficulty of information-processing in a complex world. The fundamental reasons of how suffering or mental pain is born are illustrated in Fig 13.1, which I will next go through in detail.

#### *Root causes of suffering*

The starting point is that the agent finds itself in a world which is highly complex. In such a world, acting optimally (in any reasonable sense of optimality) would require huge amounts of sufficiently detailed sensory input, together with huge capacities of computation. Unfortunately, in the real world where we live in, an agent cannot have any of those. These three root causes of suffering—*complexity of the world*, *insufficient data*, and *insufficient computation*—are shown in the left-most column of the figure. Certainly, the agent’s limited physical capabilities to act in the world and change it—catch any prey it wants, for example—create suffering as well, but here we focus on limitations related to information-processing because it can be modified more easily.<sup>1</sup>

---

<sup>1</sup>It could be argued that lack of a good model, or “inductive bias”, is another limitation. Inductive bias can refer to slightly different things: on the one hand, it is sometimes simply used as a fancy term for a Bayesian prior in a probabilistic model, but it can also refer to constraints that are more structural in the sense of, for example, the choice of the family of nonlinearities, regularization, or other computational structures used in a model (which could still, in most cases, be seen as Bayesian priors in a hierarchical Bayesian model). Basically, what we are talking about here is that the agent might not have a good model family from which to pick its model of the world, and might in particular suffering from overfitting (see footnote 4 in Chapter 4). I take here the viewpoint that, fundamentally, a good inductive bias is only necessary because the data is limited: if the data were infinite, the proper inductive bias could be learned

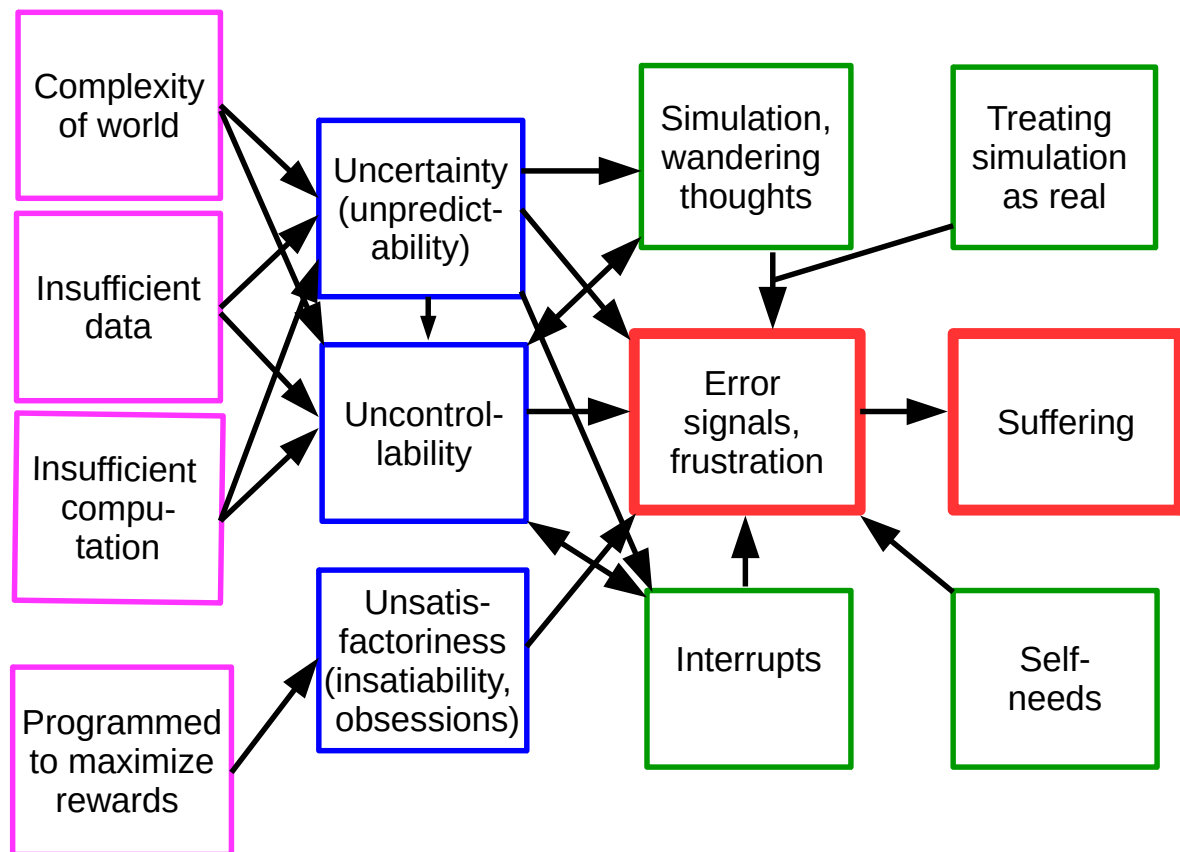


Figure 13.1: Recapitulation of the causes and mechanisms of suffering explained in earlier chapters. The boxes in magenta are intrinsic properties of the world—including the agent—while the boxes in blue are more concrete problems they pose for information-processing. The green boxes are possible functionalities in a highly developed cognitive system, and the red boxes are the postulated system finally generating suffering.

At the same time, the system is typically programmed to try to relentlessly *maximize rewards* which are ultimately set by the “programmer”—which is evolution for humans, possibly complemented by some cultural processes. The rewards are not designed to make the agent happy, but to fulfill some objectives of the programmer, such as spreading your genes in the case of evolution. It is not possible for the agent itself to decide that it wants to pursue some new objectives and re-define its own rewards; nor can it decide that it has had enough rewards and does not need more. This is a fourth root cause of suffering, shown at the bottom of the left-most column of the figure.

### *Three fundamental problems in information-processing*

The four root causes just mentioned lead to a number of challenges for information-processing, which I here condense into three: uncertainty (including unpredictability), uncontrollability, and unsatisfactoriness (consisting of insatiability together with evolutionary obsessions).

First, the overwhelming complexity of the world leads to *uncertainty*: the agent is not able to accurately understand what happens in the world. It is not even able to accurately perceive most phenomena in the world. It will try to divide the perceptual inputs into categories, but such categories are fuzzy, sometimes arbitrary, and categorization is often uncertain. Thus, there are various kinds of uncertainty.

One important special case is uncertainty about the future, also called *unpredictability*. (It is closely related to the Buddhist concept of impermanence, as will be discussed in Chapter 14.) On the one hand, it is clear that it is difficult to predict the future if even the current state of the world is uncertain, which is a consequence of uncertain perception. However, unpredictability is actually a more general phenomenon than the uncertainty of perception: even if the perceptions were very accurate and certain, it might not be possible to predict the future accurately due to the great complexity of the world. It might be impossible to learn to model the world accurately enough, or using such models might require overwhelming computational power. This is well-known in the natural sciences, where even extremely accurate measurements of a natural phenomenon do not necessarily mean you can predict it, because the prediction would require overwhelmingly advanced scientific models.

Uncertainty and unpredictability necessarily lead to *uncontrollability*, lack of control of the world: if the agent does not know what is actually happening in the world, or it does not know how to predict what will happen in the future, it cannot possibly control the world. In fact, control requires the capability to predict the results of your actions, which requires not only a good model for prediction, but also an accurate perception of the current state of the world.<sup>2</sup>

Uncontrollability is increased by other factors in addition to uncertainty. Even if the world could be perfectly perceived and predicted, there would still be uncontrollability due to at least two reasons. First, the agent has limited physical capacities to influence the world. Second, limitations of computation reduce controllability in many ways: the computational complexity of the search tree precludes finding perfect solutions to the planning problem, while the parallel and distributed nature of the agent’s cognitive system precludes proper

---

from the data by testing the performance of the models on a new test set which was not used in the learning (Feinman and Lake, 2018). Therefore, I do not discuss inductive bias in any detail in this book, and subsume the problems due to lack of correct inductive bias under the heading of “insufficient data”.

<sup>2</sup>For those conversant in Buddhist philosophy, this is similar to the idea that impermanence feeds into no-self, when impermanence is seen as related to uncertainty and no-self is interpreted as uncontrollability (Mahasi, 1996). See footnote 3 below for more on such analogues.

control of the agent's internal functioning. All these reasons make sure that the world is uncontrollable for humans. But it can be rather uncontrollable even for a thermostat, since the temperature of most environments obeys extremely complex natural laws that are beyond the understanding of the thermostat, and errors cannot be avoided.

Meanwhile, the programming of the agent to maximize rewards means that the agent finds that no amount of rewards is enough: it is *insatiable* (see page 59). In fact, the very *raison d'être* of the agent is to maximize the rewards set by its programmer or evolution. But it will never be satisfied, and the desires will never be satiated. A related property is what I called *evolutionary obsessions* (see page 59), which means that humans are compelled to seek various rewards which they might, if they thought about it rationally, prefer not to seek—such as unhealthy food and excessive competition for status. Seeking those unnecessary rewards increases the chances for frustration, thus leading to more suffering, and may even lead to less reward in the long run (by ruining your health, for example). Yet, it is difficult for humans to change what they find rewarding. I group these two properties under the umbrella term *unsatisfactoriness*, expressing the general idea that even if the world were completely known and controllable, there would still be suffering due to the fact that the system is never satiated and strives at questionable goals.

The three fundamental problems of uncertainty, uncontrollability, and unsatisfactoriness are shown as blue boxes in Figure 13.1, the second column.<sup>3</sup>

### *Error signals and suffering*

Because of uncontrollability and uncertainty, the agent will have error signals. Often, things do not go as planned or as expected, thus predictions have errors and expected reward will not be obtained, which generates error signals. Such error signals are particularly frequent because the agent is never satisfied and is always looking for more reward. Error signalling is the central red box in the third column in the figure.

Our fundamental hypothesis in this book is that such error signals are what produce the feeling, and ultimately the conscious experience, of suffering. The suffering due to error signals is especially strong if the error is frustration, i.e. the agent is trying to reach the goal (or a reward) but it fails. Thus, error signals finally lead into the red box of suffering, on the right-hand side in the figure.

<sup>3</sup>The three problems or challenges could be seen from two different viewpoints: either as properties of our natural world (at least if unsatisfactoriness is seen from a more general perspective) or as properties of information-processing in any sufficiently complex world. Here I take the latter view; the three problems are in fact created by those properties of the natural world which are given in the magenta left-hand column in the figure. These three problems are a rough analogue of what is called the three characteristics of existence in early Buddhist philosophy: impermanence (*anicca*), no-self (*anattā*), and unsatisfactoriness/suffering (*dukkha*). Impermanence is to some extent a special case of uncertainty, as will be discussed in detail in Chapter 14 on page 163. Uncontrollability is an important aspect of no-self philosophy (see page 125) and may have been its original meaning in the earliest layers of Buddhist philosophy. The Buddhist concept of *dukkha* has the broadest definition of them all, simply meaning “suffering” in one interpretation; thus our concepts of insatiability and evolutionary obsessions are only some of its aspects, as will be discussed in Chapter 14. We could have added another blue box depicting “emptiness”, a widely used concept in later Buddhist philosophy: A discussion on emptiness is postponed to Chapter 14, where it will be introduced as an umbrella term for fuzziness, subjectivity, and contextuality, and related properties. Alternatively, we could have added another box giving distributed processing and possible lack of central executive (discussed in Chapter 11) as a necessary computational consequence of the root causes on the left; now distributed processing is not explicitly mentioned in the graph, although several boxes are related to it. One more possibility would have been to introduce nonstationarity (discussed later in footnote 7 in Chapter 14) as a root cause, but it can be simply seen as a special case of the complexity of the world, even if a particularly very important one.



*Optional processes that increase suffering*

Several further processes may further be activated, depending on how sophisticated the agent is. If its cognitive architecture uses *wandering thoughts*, we get another box in the flowchart (top row, third column). It is a type of information processing that takes place only in highly sophisticated agents, which is indicated by drawing the box in green. In general, highly intelligent agents may engage in *simulation* of the world in terms of planning and replay, which in humans often happens in the form of wandering thoughts. These increase error signalling by repeating or anticipating experienced errors; this is depicted as the green simulation box on the top row feeding into error signals. Since the goal of simulation and wandering thoughts is to gain more control and reduce uncertainty on the world by better learning its dynamics, there is an arrow from the uncontrollability and uncertainty boxes to the simulation box. On the other hand, since wandering thoughts increase uncontrollability in their own way, that arrow is bidirectional.

Furthermore, the agent may react in different ways towards the contents being replayed or simulated. If the simulated contents are processed almost *as if they were real*, and the various frustrations in the simulations are processed in the same way as real frustrations, this will greatly increase suffering. Otherwise, simulated error signals might not lead to suffering. This is indicated in the flowchart as the green box on the top right-hand corner which feeds into the connection between the simulation box and the error signal box, modulating the connection between simulation and error signalling as just described.

A related design principle is *interrupts*, which are useful for handling uncertainty due to unpredictability as well as uncontrollability and are seen as one aspect of emotions as well as desire in this book. Interrupts create more frustration since by interrupting behaviour they increase uncontrollability; they also impose new goals on the agent, which obviously can lead to frustration. Interrupts may also produce specific error signals not related to frustration as they use the pain signalling pathway to grab attention. Interrupts are depicted as another green box on the bottom row, feeding into error signals.

Sufficiently developed agents have various intrinsic rewards, which may be frustrated as well. The very strongest suffering actually tends to come from the frustration of self-related goals, such as survival or the self-evaluation. Such *self-needs* create new kinds of frustration and errors, such as the agent “not being good enough” in the sense of not obtaining enough rewards on a longer time scale. This is the box at the bottom right-hand corner, again in green since it is a sophisticated module, which the programmer may include in the system or not.

This flowchart explains the conditions leading to suffering on a rather abstract level.<sup>4</sup> If we consider it from the viewpoint of interventions, e.g. practices that would decrease suffering, it clearly points out that we could reduce suffering by reducing wandering thoughts, self-needs, and other processing in the green boxes. Such ideas will be considered in detail in the next chapters. The next chapters will also consider how to deal with the

<sup>4</sup>For future research, I would like to point out that many of these boxes can be quantified, although different measures are possible, and research is needed to decide which ones are useful. Uncertainty is typically quantified by Shannon entropy as defined in information theory and already used by, e.g., Hirsh et al. (2012); Friston (2010). The complexity of the world could be the number of states (or the entropy over the typical distribution over them) in a category-based world model. The amount of data available can again be quantified by information theory: In this case, it might be the Fisher information multiplied by the number of data points between the data and the parameters of an ideal world model, i.e., how much information the data contains about the world; or the mutual information between the data and the world states. Computational resources can be quantified using flops per second or a similar measure. To quantify uncontrollability, related probabilistic computations are possible (Huys and Dayan, 2009); we might also be able to use various tools from control theory, e.g. (Liu et al., 2011).

blue boxes (uncertainty, uncontrollability, unsatisfactoriness). In the remainder of this chapter, however, we look at the process of suffering from two different angles.

## Cognitive dynamics leading to suffering

A complementary viewpoint is provided by cognitive dynamics, i.e. how the different cognitive processes work and influence each other in real-time. In some sense, this is about zooming into the part of the mechanism in Figure 13.1 that leads from the blue boxes in the second column (uncertainty, uncontrollability, unsatisfactoriness) to errors and suffering. This reveals further quantities that can be intervened on to reduce suffering.

In previous chapters, we have seen a number of steps in an information-processing procedure that translates sensory input into action decisions and possibly suffering. Such steps are recapitulated and illustrated in our second flowchart in Figure 13.2. To begin with, sensory input is received from the external world; see the upper left-hand corner of the flowchart. As the black arrow in the flowchart indicates, at the next step, the agent engages in initial sensory processing. This typically leads to recognition and categorization of objects in the world as the next step, which in our basic formalism includes recognition of the state in which the agent is. (For better visualization, the order of processing is now indicated by a single long blue arrow in the flowchart.) Recognition of the state is immediately followed by computation of the valences of the near-by objects or states. (Valence means here the prediction of the reward associated with an object, or more generally the prediction of the value of a state.) Based on those valences, a number of candidate goals are chosen, which is the process of desire in its hot, interrupting form. Next, the agent may choose one single goal and commit to it, which is also called intention. Then the agent starts planning how to reach the committed goal, possibly by some kind of tree search, and executes the plan obtained.<sup>5</sup> After finishing the execution of the plan, the agent observes the outcome, and based on it, an error such as reward loss is computed. Any such error, especially frustration, may lead to suffering. Finally, the computed error will be used in a learning process to guide future actions, which in a sense closes the loop, as indicated by the green arrow. (Only an arrow from error computation to sensory processing is drawn in the figure for simplicity, but actually the error is broadcast widely.) The flowchart shows the prototypical sequence, but in reality there is, of course, more variability.<sup>6</sup>

In the middle of the flowchart, we have “limitations of information-processing”, which influences all the steps in the processing. While that has been the main theme of the whole book, here it also provides an anal-

<sup>5</sup>To give some pointers and details on these steps: Desire was defined as a process that suggests new goals in Chapter 3. Such suggestions were seen to be possible by neural network computations in Chapter 7, resembling those performed by AlphaGo. In particular, such computations give fast approximations of valences, i.e. how rewarding near-by states are, preferably taking into account the whole future by approximating the state-values. Valence computations were further linked to the generation of interrupting (intruding) desire in the framework of the elaborated-intrusion theory of desire in Chapter 8. After the selection of one of such suggested goals, planning proceeds as already described in Chapter 3, including the idea of commitment or intention to the goal that has been selected among the possible high-value states. Note that in this framework, goals are set by the agent itself, based on predictions of rewards, which come from the outside world.

<sup>6</sup>I hope this flowchart also clarifies some conceptual differences regarding the concept of suffering, which may have been slightly confounded in previous chapters. First, I emphasize that I’m not saying that errors *are* suffering, but that errors are the direct cause for suffering. Suffering is a complex phenomenon, comparable to emotions which also have multiple aspects, including in particular a conscious experience. Conscious, subjective experience is obviously not the same thing as errors computed on the level of largely unconscious information processing. Second, nor am I saying that suffering is driving learning: It is the errors that are driving learning, after some further sophisticated computations. Thus, the errors computed lead to suffering on the one hand, and, more indirectly, to learning on the other. This is why errors, suffering, and learning are separate items in the flowchart.

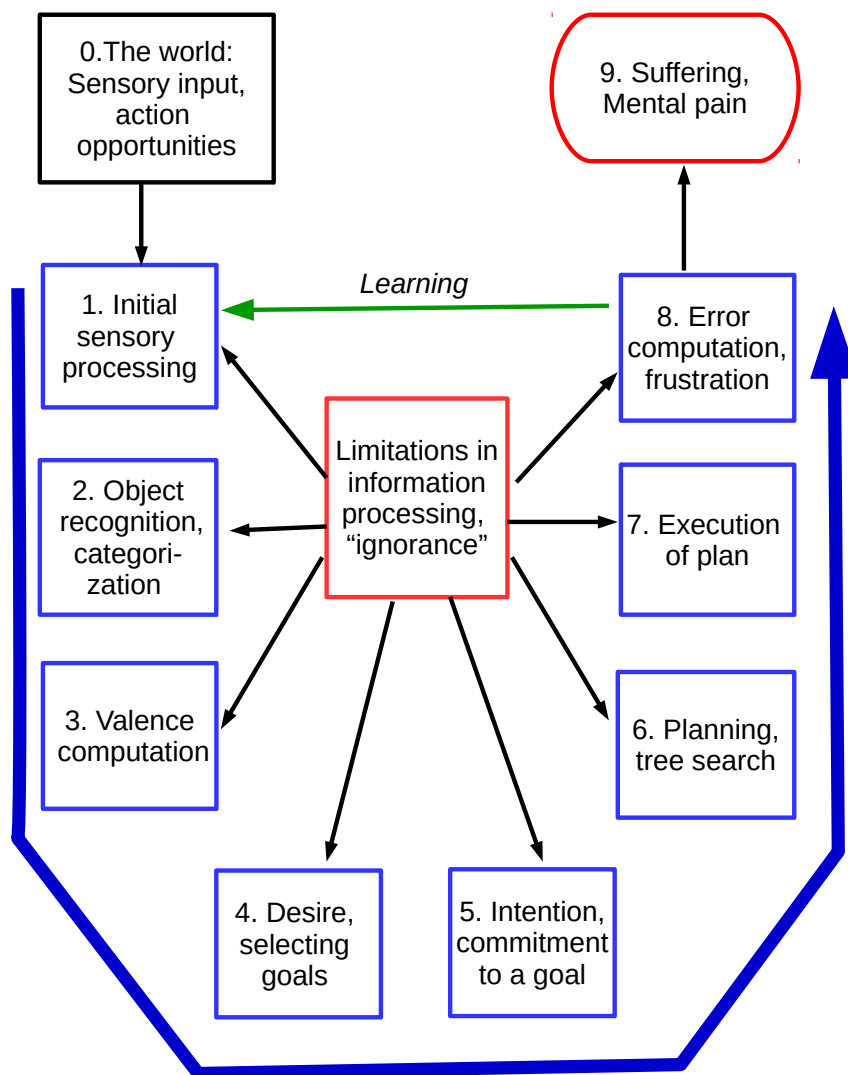


Figure 13.2: Recapitulation of the detailed cognitive mechanism of suffering explained in earlier chapters. The sensory input from the world enters the system in the top left-hand corner of the chart (step #0). It is processed in a sequence of steps #1-#7, along the big blue arrow. Throughout these steps, there are various limitations in the information-processing, or “ignorance”, to use a Buddhist term defined in the main text. Finally, the computation leads to computation of errors (step #8) which may lead to suffering in the top right-hand corner of the chart (step #9). The error computations are further fed back to the whole system as indicated, for brevity, by the single green arrow closing the cycle and going back to sensory processing.

ogy to Buddhist philosophy which uses the term *ignorance* in connection with similar conceptual schemes. Ignorance describes the fundamental underlying reason why the agent's cognitive apparatus creates so much suffering—thus adding to the adverse properties of the external world, as shown in the previous flowchart. We have already seen various kinds of information processing which might be called ignorant in the sense that they can be seen as faulty or lacking, and they increase suffering. Thus, limitations of information processing could be seen as one definition of “ignorance”. However, this Buddhist term could also be interpreted more literally in the sense of our being ignorant of those limitations; a “meta-ignorance” so to say. We can thus consider such forms of ignorance as 1) ignorance of the arbitrariness and even harmfulness of our rewards and goals, 2) ignorance of the uncertainty and fuzziness of perceptions and concepts, 3) ignorance of uncontrollability, and finally 4) ignorance that the simulation is not real—but this list is not meant to be exhaustive.<sup>7</sup> In all these cases, we can claim, inspired by Buddhist philosophers, that there is some kind of a “mistake” in our ordinary thinking and functioning of the brain. In particular, the mistake, or flaw, is about misunderstanding where suffering comes from and how it can be reduced. Importantly, in contrast to fundamental limitations of information processing, this “meta-ignorance” could be corrected. Most of this book has actually been devoted to explaining what such ignorance consists of and how its different forms lead to suffering or amplify it. So, in Figure 13.2, ignorance is naturally placed in the very middle.<sup>8</sup>

<sup>7</sup>In early Buddhist philosophy, a central form of ignorance is the belief in “self”. I omitted it from this list because in this book, I have attempted to largely reduce such no-self ideas to less abstract concepts, in particular uncontrollability.

<sup>8</sup>It is interesting to compare the mechanisms described above with a central idea in Buddhist philosophy: the twelve-link chain, which has served as a central inspiration for this flowchart. (A related approach is proposed by Grabovac et al. (2011).) The Buddha elaborated his idea of desire as the origin of suffering by building a sequential causal model, which can be seen as an instance of his more general ideas on “causality” resumed under the heading of dependent co-arising/origination (Mahasi, 1999; Anālayo, 2003). While different variants exist, I consider here the version with twelve items. Some of the items in his chain correspond clearly to concepts we have seen in this book, while other do not. The chain begins by three items which are difficult to interpret, and seem to be metaphysical speculation about how the ignorance of the true nature of reality creates consciousness and this creates the world. In the text above, I provided some more scientific interpretations of “ignorance” in terms of limitations of information-processing, and our ignorance of those limitations and their implications. After those initial three items, the middle part of the Buddha's chain goes as follows: 4) “Name-and-form”. This basically means the world, including our internal world of memories and consciousness. 5) “Six-fold sense bases”. This is when the sensory organs receive input from the outer world, or memories or wandering thoughts enter the mind (which is in Buddhism considered a sixth sense). 6) “Contact”. This I interpret as perceptual processing leading to object recognition, where the brain processes information and interprets the incoming stimulus in terms of a given category (“That's a dog”). 7) “Feeling” (*vedanā*) means computation and perception of the valence of the sensory stimulus: Is it good or bad, do I like it or not? 8) “Craving (or thirst, or desire)” is the same as desire in our terminology, as always including aversion. You may want the object you has seen, or you may want to get rid of it. A number of goals are considered at this stage. 9) “Attachment (or clinging)”. I proposed at the end of Chapter 7 to interpret this as forming an intention, i.e. committing to a certain plan and a goal, and planning for it. (In Buddhist literature, the interpretation of “attachment” is actually highly variable, and in my view rather muddled: often, the distinction between desire and attachment is only a matter of degree. It is sometimes pointed out that an alternative translation of the word in question (*upādāna*) is “fuel”, which might give an alternative interpretation as being related to the learning process in the next steps.) 10–11) “Becoming and birth” are the next two steps which are a bit more difficult to interpret, and have traditionally been interpreted in more metaphysical terms. I suggest interpreting them as referring to the *learning process* which creates various associations in the mind, including creating habits out of one's actions. Thus, whatever the agent does leads to “birth” of new action tendencies and associations. 12) “Old age and death” is the final result of the above causal chain, and can be interpreted as simply “suffering”.

We see that steps 4–9 here directly correspond to the first boxes in our flowchart (from box #0 to box # 5). The boxes #6–8 following that, including actual planning, plan execution, and error computation may be missing in the Buddhist chain, or they could be seen as being included in its steps 10–11. The steps 10–11 in the Buddhist chain are, in my tentative interpretation, specifically related to the ensuing learning process, shown by the horizontal green arrow in the flowchart. Step 12 is a poetic description of suffering in the red box #9 at the top right-hand corner in the flowchart. In any case, the two models share the crucial idea of how sensory input leads to

It should be noted that this series of processing steps is not only started by external stimuli (the sight of something good) but also by internal simulation, which is another reason why the processing constitutes something more like a cycle or a loop. For example, wandering thoughts, or almost any kind of thoughts, trigger memories or predictions of sensory stimuli, and the cycle is launched almost as if those stimuli were real. This is, however, not shown in the flowchart for the sake of brevity.

## An equation to compute frustration

While the flowcharts above help us understand the mechanisms behind suffering and even design interventions, the real strength of a computational approach as taken in this book is that we can quantify things, at least in principle. I don't mean that we would necessarily be able to give a number measuring the strength of suffering, but we can understand the connections between different quantities more explicitly than with flowcharts. As the most powerful recapitulation of the theories of preceding chapters, I next propose a simple equation that describes the amount of suffering experienced by the agent. This will be the central approach in the next chapters, where we attempt to reduce suffering.

The starting point for the equation is Chapter 5 (page 52), where we defined suffering as reward loss, that is, the difference between the expected reward and the obtained reward. (More generally, the reward prediction error might be used.) This theory provides a quantitative basis for modelling suffering. Reward loss is based on a simple mathematical formula, so we can look at the different terms it contains. We can analyse how they influence reward loss, and how they could eventually be manipulated. The equation we had in Chapter 5 was, however, very basic and did not take into account any of the complexities of a real cognitive system that we have seen in later chapters. So, we need to look at the different factors influencing reward loss in more detail and reformulate the equation.

The first point to consider is that any quantities affecting reward loss need to be perceived by the agent. While the difference between expected reward and the reward actually obtained is, in principle, the basis of suffering, the difference cannot of course cause suffering by itself. It has to be computed—that is, perceived—by the agent. So, we need to make a connection between limitations in perception and categorization on the one hand, and frustration on the other. Due to such limitations, the perception of the actual reward is uncertain and subjective, as explained in Chapter 10. Obviously, our expectations are subjective and may be overblown as well. Yet, the agent can only compute the reward loss based on its own perceptions, on the information at its disposal.

The second point is that as with any perception, the level of certainty of the error computation should also be taken into account: If the perception of reward loss is particularly uncertain (say, because it is dark and you cannot see what you get), this should reduce the effect of reward loss. It is common sense that if the agent is not at all certain about what happened, it should not make any strong or far-reaching conclusions, and the error signal should not be strong.

Furthermore, again following general rules governing perception outlined in Chapter 10, the intensity of the perception of reward loss is modulated by the attention paid to it. Reward loss causes less suffering if little attention is paid to it, for example, when one is distracted by something else—simply because you might not even notice reward loss occurring. Paying attention to something may also be necessary for becoming

---

desire, which via attachment (clinging/intention) leads to suffering, and how this suffering somehow perpetuates itself by a learning process.

conscious of it.<sup>9</sup> Thus, the amount of attention paid to the reward loss must be included in the equation. There are also further related phenomena which change the amount of frustration experienced: we may take the contents of such simulation more or less seriously (Chapter 12), and we may find errors acceptable if we are deliberately trying to learn something new. For simplicity, we include these aspects in the term called “amount of attention” being paid to the reward loss, since not taking simulation seriously is related to not paying a lot of attention to it.

The final important point is that error computation can happen many times, in particular in the case of replay or planning, which means we perceive, or rather simulate, the same (possibly imaginary) reward loss again and again (Chapter 9). If you replay an event just once in your head, the suffering is multiplied by two, almost.

Taking all these aspects into account, we arrive at the following which I call the *frustration equation*:

$$\begin{aligned} \text{frustration} = & \\ & \text{perception of (the difference of expected and obtained reward)} \\ & \times \text{level of certainty attributed to that perception} \\ & \times \text{amount of attention paid to it} \\ & \times \text{how many times it is perceived or simulated} \end{aligned}$$

In this equation, we have four terms on the right-hand side, i.e. after the equality sign, multiplied by each other. First, there is the basic formula of reward loss in parentheses. Thus it includes the amount of expected reward and the amount of obtained reward, whose difference is computed. As with reward loss, if this difference is negative, it is set to zero—if the obtained reward is greater than the expected, there is zero frustration. But crucially, the reward loss here is modulated based on how it is perceived by the agent.<sup>10</sup>

Then, this perceived reward loss is multiplied by three modulating factors. We use here multiplication to emphasize how the perceived reward loss may actually lead to no suffering at all if just one of these modulating factors is zero. The modulating factors are: the level of certainty that the agent attributes to the perception of reward loss (zero meaning absolute uncertainty, one meaning complete certainty), the amount of attention paid to reward loss (including how seriously the contents of the consciousness are taken), and finally the number of times the event is perceived (taking account of the fact that it may be replayed or simulated in planning so that it is “perceived” more than once).<sup>11</sup>

It should be emphasized that such frustration happens on different time scales: from milliseconds to even years, perhaps. In the smallest timescales, the suffering is likely to be much weaker than on longer time scales; see Chapter 7 (page 85) for discussion on this topic. The equation above is intended to be applied separately

<sup>9</sup>See footnote 18 in Chapter 12 for discussion on this connection.

<sup>10</sup>There is a subtle point about perception of the reward loss, which is that the system may first compute percepts of the two quantities (expected and actually obtained rewards) and then compute the difference, or it can directly attempt to compute the percept of the difference. In other words, we can have perception of the difference, or the difference of the perceptions. (In the case of the expectation, its “perception” might rather be called the “estimation” of the expectation.) Intelligent systems may use either of these two approaches. I shall not venture into speculating which might be the case in the human brain.

<sup>11</sup>The equation does not mention the “self”. This is because I reduce self-based suffering to frustration of self-needs based on the logic explained in Chapter 6; that is the logic used in the following chapters. Alternatively, inspired by Buddhist philosophy, we might think about adding another multiplying factor to the equation, called “relevance to self”, which would measure if the reward loss is affecting the self (in some sense to be defined).

on different time scales.<sup>12</sup>

Another point that is useful to recall here is how we reformulated action selection as based on rewards only in Chapter 5. While we often talk about goals and planning, for example in the flowchart in Figure 13.2, the goals are now seen as something that the agent itself sets in order to maximize rewards. In particular, Chapters 7 and 8 explained how an agent would predict that a certain state gives a big reward, then set it as a goal state, and start planning for it—the same logic underlies Figure 13.2. Thus, goals are not something inherent in the world, but rather a computational device used by the agent in order to maximize rewards. That is why this equation uses the formalism of reward loss, instead of the basic formalism of frustration of goals initially used in Chapter 3. Still, goals are implicitly present in this equation since the expected reward is often the reward that reaching a certain goal would give (according to the agent's prediction).

Now, the essential point here is that all the terms on the right-hand side of the equation are something that can be influenced or intervened on, at least to some extent. It is possible to develop methods that change the terms on the right-hand side, thus changing the amount of frustration and suffering. The next two chapters are largely an explanation and application of this equation from that viewpoint.

In fact, so far, it might actually seem that the book has been just one big complaint. Suffering seems unavoidable, a necessary consequence of intelligent information processing. However, in the following final chapters, we will see a way forward: what an intelligent agent can do to actually reduce its suffering.

---

<sup>12</sup>It should be useful to formulate a similar equation based on RPE instead of reward loss. Such a formulation would handle these complex temporal aspects in a more principled way, and in particular, it would encompass frustration based on expectations alone, as treated in Chapter 5, especially footnote 20. I leave that for future research.

## Chapter 14

# Reprogramming the brain to reduce suffering

In this chapter and the next, I will present various ideas on how to reduce suffering in a complex intelligent system acting in a complex world—such as humans. I derive various ways how information-processing should be changed, i.e. how the agents should be reprogrammed, based on the theory presented in this book. Since the systems in question, such as our brain, have largely learned their function from input data, an important part of such reprogramming is retraining the learning system by inputting new data into it.

The methods discussed here are not original at all: almost all come from Buddhist and Stoic philosophy or related systems. The goal here is to interpret them from a computational AI perspective, using the theory developed in this book. Thus, we gain more understanding on how they work, why they work, and what could be done to improve them.

The main starting point in this chapter is the frustration equation we just encountered (page 158). We can try to reduce suffering by changing any of the terms on the right-hand side of the equation, since that inevitably implies that the frustration on the left-hand side of the equation is reduced. We can see from the equation that the obtained reward should be increased, because it has a negative sign in the difference computed. In contrast, all the other terms on the right-hand side should be decreased because their contribution to frustration is positive.

Maximizing the obtained reward is really a very conventional way to try to reduce suffering, based on the wide-spread view that happiness comes from having achieved all your goals, and having got what you wanted.<sup>1</sup> However, that is difficult for reasons which are rather obvious. Many resources are limited: not everybody can have the best cars, the best wines, and the best sex partners. There is fierce competition over such resources, and not everybody can win. Besides, expectations are adapted to the obtained level of rewards, so what used to feel good no longer brings happiness after a while, as discussed in Chapter 5 (page 59).

So, instead, we attempt here to *reduce all the terms other than obtained reward* in the frustration equation. In this chapter, we consider how it is possible to reduce two of those terms: the (perception of) expected reward, and the certainty attributed to the perception of reward loss. (The next chapter will consider reducing the remaining terms, as well as some further methods.) Such reduction also includes reducing self-needs as a special case, thus complementing the frustration equation by the logic of the flowchart in Fig. 13.1. Ultimately, such practices lead to reducing all desires and aversions. This approach may be rather unusual in the context

---

<sup>1</sup>(Oatley and Johnson-Laird, 1987; Van Boven, 2005; Heathwood, 2015); cf. Diotima in Plato's *Symposium*: “[T]he happy are made happy by the acquisition of good things”. For a discussion of different definitions of “well-being”, which can be seen as equivalent to happiness here, see Crisp (2017); Fletcher (2015).



of modern Western psychology and philosophy, but it is thoroughly standard in Buddhist and Stoic philosophy.

## Reducing expectation of rewards

Let us first look at the term “perception of difference of expected and obtained reward”, i.e. perception of reward loss, in the frustration equation. This should be made as small as possible, ideally zero or even negative. As already mentioned, the most conventional way to reduce it would be to try to increase actually obtained rewards, but that is very difficult. So, we need to do something more clever. A well-known idea in Buddhist and Stoic philosophy is to lower your expectations—at least in the colloquial sense of the word. Then, your reward loss should be smaller, and will perhaps vanish altogether.

The expected reward is essentially a product of two things: the probability you assign to obtaining the reward, and the actual amount of the reward if it is obtained (considering the basic case where the amount of reward, if obtained, is fixed). Thus, reducing the expectation of a reward can be accomplished by either reducing the probability the agent assigns to it, or the value it sees in the reward. This can be compared to a lottery. Suppose your initial chance of winning a Porsche is 1%. Obviously, the lottery would be made less attractive if the probability of winning is lowered to 0.01%; it would also be less attractive if you realized that the Porsche is second-hand and not so cool after all. In both cases, your expected reward is reduced.

Most importantly, rewards in the real world are always a bit subjective, and so are the probabilities we assign to them. A new Porsche may feel like a great reward to one person, while it may matter very little to another; this is why we have to talk about *perceived* reward. People will also have very different guesses of the probability of winning it. Since these quantities are subjective, it is possible to change our estimates of them by changing our beliefs, perceptions, and associations, even if the actual physical reality remains unchanged.<sup>2</sup>

A key goal of Buddhist and Stoic systems is exactly such re-evaluation of the probabilities and rewards. To accomplish this, Buddhist philosophy talks about the “three characteristics of existence”, which are impermanence, no-self, and unsatisfactoriness. They map roughly to our concepts of uncertainty, uncontrollability, and unsatisfactoriness we discussed in the preceding chapter.<sup>3</sup> Each of these characteristics gives a reason why the rewards are actually lower than what they would otherwise be, or what they appear to be, as will be explained next.

### *Facing uncontrollability*

Uncontrollability, discussed in Chapter 11, is a key concept here. The level of controllability is clearly related to the level of expected reward. If you think the world can be controlled, you will expect to achieve high rewards, because you think you are able to take courses of action that give you the very highest rewards, and you are

---

<sup>2</sup>The influence of the subjectivity and contextuality of perception, and ensuing possibilities of reducing frustration by changing the perception, are actually quite complex phenomena and create many further possibilities of reducing the term being considered here. Logically, we might also try to *increase* the perceived *obtained* reward independently of the actual reward obtained. This may be possible by somehow learning to better appreciate the rewards obtained, but I will not develop that idea any further here. In fact, it may not always be possible to distinguish between the actual reward and the perceived reward: the reward may not have any objective definition, as may be obvious in the case of winning a Porsche. Likewise, it might be possible to somehow reduce the perceived reward *loss* even if the actual reward and the expected reward are unchanged, but I cannot very easily think of a method to achieve that; seeing the reward loss as a useful learning signal would be a bit in that direction.

<sup>3</sup>See footnote 3 in Chapter 13 for a detailed discussion

reasonably certain that you can achieve them. In contrast, if you think the world is *uncontrollable*, you assign a low probability to achieving any rewards, and the higher rewards may seem to be completely out of your reach. Then, your expectation of reward is smaller, and you are less likely to suffer from a reward loss, i.e. frustration. This is how considering the world to be uncontrollable reduces suffering.

To transform this logic into a practical method for reducing suffering, the trick is to acknowledge the fact that you have little control and there can never be very much control, and firmly believe in that fact. We saw earlier how the Stoic philosopher Epictetus emphasizes how little we can control (page 125). He continues by explaining that if we are mistaken about this point, suffering is inevitable:<sup>4</sup>

The things in our control are by nature free, unrestrained, unhindered; but those not in our control are weak, slavish, restrained, belonging to others. Remember, then, that if you suppose that things which are slavish by nature are also free, and that what belongs to others is your own, then you will be hindered. You will lament, you will be disturbed, and you will find fault both with gods and men.

Likewise in Buddhist philosophy, the original form of the no-self philosophy says that nothing is part of me, which is a way of saying that nothing can be controlled, as we saw in Chapter 11. Understanding this is crucial according to the Buddha:<sup>5</sup>

All [mental phenomena], whether past, future or present, internal or external, gross or fine, inferior or superior, far or near, should be seen with one's own knowledge, as they truly are, thus: 'This is not mine, this I am not, this is not my self.' (...) [S]eeing thus, [the disciple] grows wearied of form, wearied of feeling, wearied of perception, wearied of volitional formation, wearied of consciousness. Being wearied, he becomes passion-free (...), he is emancipated [from processes leading to suffering].

Here, I interpret “growing wearied” as signifying that the reward expectations are lowered, or little enjoyment anticipated. Thus, the point is that recognizing uncontrollability, or inexistence of self, reduces expectations of reward, which reduces suffering.

Buddhist philosophy emphasizes the importance of understanding “causality”. Such causality means that events in the world just follow from each other based on natural laws, for example those depicted in Figure 13.2. This thinking minimizes the importance of free will and the control that the agent can have over the world; it is related to what is called determinism in Western philosophy. I would think, therefore, that the Buddhist emphasis on what they call causality is just another viewpoint on uncontrollability; seeing such causality is one way of realizing that the world is uncontrollable. Stoic philosophers advocated the study of the natural sciences (which they simply called “physics”), with a similar goal.<sup>6</sup> Such a philosophy could be criticized on the grounds that it might lead to total inaction; this point will be discussed at the end of this chapter.

---

<sup>4</sup>Paragraph 1 of *The Enchiridion*.

<sup>5</sup>*Samyutta Nikaya* 22.59, latter half, strongly shortened, based on the translation by Mahasi (1996).

<sup>6</sup>Causality is a topic of great current interest in AI (Pearl, 2009; Peters et al., 2017; Gershman, 2017). However, the meaning of the term in AI is a bit different, and in particular, very specific: It is more about the difference between correlation and causality, and how an agent could learn that difference. In AI, understanding such causality will enable the agent to act more efficiently, increase its control of the world and the rewards it obtains, as well as better predict the rewards. In contrast, in Buddhism, understanding causality is about admitting the determinism of the world and minimizing the control of the agent and free will. Eventually, both these two kinds of “understanding causality” may reduce suffering in their own ways. Briefly, the AI kind of understanding means that you can find the optimal actions, while the Buddhist understanding means appreciating how little reward even those optimal actions bring.

*Facing uncertainty, unpredictability, and impermanence*

Uncontrollability is closely related to the concept of uncertainty. Uncertainty feeds into uncontrollability: if the workings of the different objects in the world are uncertain, even quite random, the world cannot be very well controlled. Likewise, uncontrollability often leads to uncertainty of whether rewards will be obtained. So, in some sense, these are two sides of the same coin.

Buddhist philosophy focuses on the related concept of impermanence which can be largely seen as a special case of uncertainty. Impermanence means that the world is constantly changing, and usually in unpredictable ways.<sup>7</sup> For example, any object that you may possess can break or get lost. Any enjoyment that you get is likely to be fleeting. In fact, even your feelings and opinions are impermanent: today you like one thing, but perhaps tomorrow you're already bored with it and want something else; what you consider important today may have no significance to you next month. Obviously, impermanence thus interpreted leads to uncertainty.<sup>8</sup>

Going back to our frustration equation, the consequences of uncertainty are also very similar to the consequences of uncontrollability. The central point is that any future rewards are uncertain, i.e. unpredictable. Rewards and the circumstances leading to rewards can change, so an agent cannot really know whether it will get any reward after executing its plan. Thus, the agent should lower the probability it assigns to any future reward. If the agent acknowledges such uncertainty of the world, its expectations regarding rewards will be lowered, just like in the case of uncontrollability. Consequently, frustration will be reduced. (Later, I will talk about perceptual uncertainty, which has a different effect on suffering.)

It is quite paradoxical that Buddhist practice, which turns your attention to uncertainty and uncontrollability, tends to reduce stress and suffering. In Chapter 6 we saw that uncertainty and uncontrollability are usually thought to lead to more stress, not less. I think the paradox has a lot to do with one's attitude to uncertainty and uncontrollability. Somehow Buddhist philosophy seems to result in a particularly appropriate attitude, related to their acceptance which will be considered in more detail in the next chapter.<sup>9</sup>

---

<sup>7</sup>From an alternative probability theory viewpoint, impermanence could also be seen as incorporating *nonstationarity*, which is the technical term for the situation where the world is changing, and a statistical model learned on data in the past may not be valid anymore in the present and even less in the future. Such problems were already alluded to in previous chapters when it was pointed out that humans may be evolutionarily adapted to the African savannah instead of the modern city environment (Chapter 5); or that emotional reactions learned as a child may be far from optimal in an adult since the environment is very different (Chapter 8). It could be argued that nonstationarity is as important as scarcity of the data, and an independent root cause of suffering. However, I don't pursue that line of argumentation here since I tend to think that any problems with nonstationarity could be avoided if the agent had enough data and computation since it would then be able to predict the nonstationarity (like Laplace's demon in footnote 15 below); but this is admittedly a complex and controversial point that needs further research.

<sup>8</sup>However, the Buddhist impermanence has also aspects that cannot be considered to be forms of uncertainty. For example, if you know for sure that you will die tomorrow, there is no uncertainty, although this is the quintessential expression of impermanence in Buddhism. In fact, in a meta-level sense, the central point in early Buddhist philosophy of impermanence is that impermanence itself is absolutely certain, as well as "permanent": Everything will perish one day. Some exceptions may exist, however: "enlightenment", nirvāna or nibbāna, is permanent according to most schools (Harvey, 2009, p. 52), consciousness is permanent for some Yogācāra thinkers (Williams, 2008b, p. 99), and a rather obscure metaphysical construct called *dharmakāya* is also permanent in some Mahayana schools (Williams, 2008b, p. 106).)

<sup>9</sup>This contradiction between ancient philosophers and modern psychology has baffled many commentators. I would speculate that the problem is that not all research on lack of control (or uncertainty) has clearly made the distinction between the level of control/certainty the agent *perceives* to have, and the level of control/certainty the agent *wants* to have. If the agent wants to have more control or certainty than it perceives to have, this leads to a case of frustration, on a "meta-level" or of self-needs, and thus suffering. A Buddhist practitioner is supposed to accept that control is not possible, and everything is uncertain. Thus, she gives up wanting any control and avoids any such meta-level frustration. Combined with the fact that frustration (on the ordinary level) is reduced, as

*Facing unsatisfactoriness*

In Buddhist philosophy, the two characteristics of impermanence and no-self (roughly, uncertainty and uncontrollability) imply a third characteristic: unsatisfactoriness, which has many meanings and interpretations. On the one hand, it expresses the idea that whatever we try to achieve, we often fail due to uncontrollability and uncertainty. In this sense, it simply recapitulates the aforementioned properties. On the other hand, unsatisfactoriness can be seen as an extremely general characteristic, which is, in Buddhist philosophy, assumed to penetrate all phenomena and existence. In fact, in the original Indian texts, the single word *dukkha* is used to express such unsatisfactoriness as well as suffering, i.e. this very thing we are trying to reduce. One could express the relation between these two meanings by saying that all phenomena are unsatisfactory in the sense that they can produce suffering, one way or another.<sup>10</sup>

In Buddhist philosophy, it is recommended to acknowledge the unsatisfactoriness of all phenomena. This can be justified using our frustration equation: if the agent is strongly convinced about the unsatisfactoriness of all phenomena, its expectation of reward will be very low, and thus reward loss will be small and rarely even occurs. Thus, here we are talking about a very general, if a bit vague, strategy for lowering the expectations of rewards. Even in cases where it is not quite obvious to see how uncertainty or uncontrollability apply—perhaps you can get chocolate quite easily and there is little uncertainty about that—it is still possible to think that the phenomena concerned are unsatisfactory, for example, because there are various negative side effects hidden in them (more on this below).

As a training method of great generality, the Stoics suggested reviewing any plan of future action with the view of anticipating what could go wrong and how the plan will not lead to great enjoyment after all. Epictetus gives a famous example of going to a Roman bath:<sup>11</sup>

If you are going to bathe, picture to yourself the things which usually happen in the bath: some people splash the water, some push, some use abusive language, and others steal.

With this mindset, you will not expect much enjoyment, i.e., reward, and you will not be disappointed. Such a scenario could be analysed in terms of uncontrollability and unsatisfactoriness as well, but unsatisfactoriness may be a more natural viewpoint.<sup>12</sup>

In Chapter 13, we defined unsatisfactoriness in a more specific way. We complemented the properties of uncontrollability and uncertainty by two phenomena grouped under the title of “unsatisfactoriness”. First, we had *insatiability*: Chapter 5 discussed at length the idea that an intelligent system which is programmed to maximize reward will never be satiated or satisfied, by the very construction of the system. It will never find that it has had enough; in a word, the system is infinitely greedy. This points at one reason why simply getting a lot of rewards will not remove reward loss in a sufficiently intelligent agent: in the long run, getting more reward will increase the expectation of rewards. The second aspect of unsatisfactoriness in our framework was

---

argued in this chapter, this should lead to less suffering.

<sup>10</sup>On the translation of *dukkha*, see Anālayo (2003, p. 244).

<sup>11</sup>Paragraph 4 of *The Enchiridion*.

<sup>12</sup>In particular, such mental imagery may serve to increase the perceived probability of adverse events based on what is called the availability heuristic (Tversky and Kahneman, 1974). That is, humans tend to estimate the probability of events based on how easily they can recall (or imagine) those events. If you willfully imagine an event happening in the future, that will make the event more accessible in terms of memory retrieval, and thus you may start considering its probability of happening is higher. When things going wrong are perceived to have a higher probability, expected reward is necessarily reduced. Thus, imagining adverse events may be a particularly powerful way of influencing unconscious decision-making processes.

*evolutionary obsessions* (Chapter 5). Even the very goals pursued and the rewards obtained can be questioned. Perhaps the evolutionary system gives you a certain reward for drinking a sugary drink. But we know very well that such a reward is misleading: The drink is not good for you when all its effects are considered in the long run. In these two senses, rewards are often deeply unsatisfactory.

From this viewpoint of unsatisfactoriness, even if we could totally control the world and everything were certain, the result of our strivings would not be that great anyway because it would not produce a lasting satisfaction or pleasure. While uncertainty and uncontrollability are more about the probability of getting various kinds of rewards, unsatisfactoriness (both in our sense and the Buddhist sense) is really about the worth of the rewarding objects or events themselves, once the agent has obtained them. Even the very best chocolate, if you eat it every day, will ultimately leave you indifferent, and may ruin your health in the long run.

This latter logic of unsatisfactoriness actually works a bit outside of the frustration equation because it is not that the rewards or their probabilities (or any other terms in that equation) are changed: it is rather understood that even if the rewards are obtained, there are unpleasant side effects. The frustration equation is in a sense short-sighted: it only considers the direct, immediate effects of rewards or their simulation. In contrast, the ideas of insatiability and evolutionary obsessions bring a longer time scale into the picture, pointing out that obtaining rewards now may actually increase frustration and suffering *in the long run*. In particular, this logic of unsatisfactoriness reduces desires, so there is less opportunity for any frustration to arise, as will be considered in detail later in this chapter.

### *Frustration due to aversion*

It may be easy to see how frustration is reduced by lowering any expectations of enjoyment, say when going to a public bath with Epictetus. However, it may be more difficult to see why the aforementioned attitudes would also reduce frustration due to aversion.

Let us consider aversion based on expecting that a bad thing is likely to happen, such as your neighbours starting a noisy renovation. Now, taking account of uncontrollability means you cannot really avoid the bad thing, at least not with any certainty. This means that the probability of the bad thing happening is larger than what you might have initially thought—your flat will be noisy for sure. Thus the expected reward is less than what you would have thought without taking uncontrollability into account. More precisely, it is more negative, since the probability of a negative reward is larger. Thus, frustration is reduced by reducing the expectation of reward by making the negative expectation even more negative.<sup>13</sup>

Likewise, thinking in terms of unsatisfactoriness (in the Buddhist sense) means thinking that the bad thing is likely to be really bad—the noise is probably going to be something quite unbelievable. Again, this reduces expected reward in the sense of making it even more negative, and what actually happens is less likely to give you a negative surprise and frustration. Thus, admitting both of those two characteristics, uncontrollability and unsatisfactoriness, lead to reduction of frustration.

<sup>13</sup>Here, I'm assuming a kind of complete symmetry between desire and aversion: they produce suffering by the same mechanism captured by the frustration equation. That is, aversion is just the same as desire, the main difference simply being that the expected reward is negative, while the obtained reward could be even more negative. It could be argued that this is not the case because aversion *in itself* produces suffering independently of any frustration. If I feel disgust or fear, surely I am suffering from those emotions alone. I have actually argued earlier that interrupts create suffering in themselves because they use the pain signalling pathway to grab attention (see page 95). However, I have also argued that desire in itself creates suffering, just like those negative emotions (see page 86), so this may not break the symmetry between (positive) desire and aversion. This is a complex question that I leave for future research.

A classical Buddhist account would further point out that impermanence means that the object of aversion will eventually disappear, which makes at least the feeling of aversion weaker. Clearly, it will give me some comfort knowing that the noise will not be there forever. This brings the property of uncertainty, of which impermanence is a special case, into the framework of frustration, although in a bit indirect way.<sup>14</sup>

## Reducing certainty attributed to perception and concepts

Another term in the frustration equation that we can reduce is the certainty attributed to the perception. As we saw in Chapter 10, perception is uncertain. To recapitulate the main ideas: perception is based on limited data, thus necessitating unconscious inference, which may not always be much better than guessing. Perception is also subjective: different people can have different priors and thus different perceptions. Subjectivity is made even more serious by the strong selection of incoming information by attentional mechanisms. Since the computational capacity is always fundamentally limited, and the world is awesomely complex, it is not possible to build a perceptual system that always makes correct inferences, let alone one that perceives the “true” reality. This should imply a fundamentally skeptical attitude towards any perception: we should not make too strong conclusions based on sensory input.

Thus, we see that uncertainty has two different aspects. There is the objective unpredictability of the world: surprising and unexpected things can happen, the world is to some extent random—this is the kind of uncertainty we focused on earlier in this chapter when talking about impermanence. But here, we focus on the uncertainty in our perceptions and beliefs of the world, which I call here perceptual uncertainty. The point is that we don’t know with any great certainty what the state of the world is, since we have neither enough data nor enough computation to perceive it properly.<sup>15</sup> Such perceptual uncertainty increases the effects of unpredictability and uncertainty that we saw earlier, since it makes the world even more unpredictable for the agent.

There is also a completely new aspect to perceptual uncertainty: it is relevant when evaluating the reward loss or frustration. After the agent has completed an action sequence in view of getting reward, it tries to evaluate the reward loss. Now, if the agent is wise enough, it will understand that it cannot know with certainty how much reward it got. A drink may have tasted good, but you don’t know if it was actually good for you. That is why in the definition of reward loss, we should really be talking about *perceptions* of rewards instead of any objective quantities; this is precisely what is done in the frustration equation.<sup>16</sup> Since the reward loss is

<sup>14</sup>Putting uncertainty into the framework of the frustration equation is not straightforward in the case of aversion, as seen in this example. Taking account of the various forms of uncertainty might mean that you realize that the bad thing, which you initially thought is certain, is actually less likely to happen than what you first estimated. Paradoxically, this increases the expected reward, because the negative reward is less likely to happen, and actually increases your frustration—that is why the effect of admitting uncertainty on aversion-based frustration is not straightforward. See also the discussion of the connection between fear and frustration in Chapter 6; the current example is clearly one of fear since it is about expecting something bad to happen.

<sup>15</sup>This division into two kinds of uncertainty could be criticized on philosophical grounds. Laplace proposed that an intellect (called a “demon” by later commentators) which knows everything about the world, would be able to perfectly accurately predict everything, and nothing would be uncertain to it. Thus, from this viewpoint, uncertainty is always a reflection of ignorance about some aspects of the world. I shall not go into that debate here, and just acknowledge that the division I make here is not very rigorous, while in line with how randomness is handled in AI theory, and hopefully also in line with the common-sense idea of randomness. It is closely related to the distinction between “epistemic” and “aleatoric” uncertainty discussed by Hüllermeier and Waegeman (2021).

<sup>16</sup>To emphasize: our basic definition of reward loss on page 52 does not take into account the fact that it is perceptions that matter. Obviously, it cannot then take the uncertainties into account either. Thus, the definition must be changed accordingly, and this

uncertain, any conclusion drawn from it should not be given too much weight, according to the basic principles of Bayesian inference.<sup>17</sup> Many philosophers over the centuries have pointed out that what first appears to be a negative outcome may even turn out to be positive, and vice versa.<sup>18</sup> Furthermore, since thoughts are derived from perceptions, the agent should be skeptical towards its own thoughts as well.

If the agent is programmed to take account of the fact that all its perceptions are uncertain, it would likely have weaker reward loss signals. Consider an agent that attempts to get some chocolate. Suppose that after executing a plan, the agent is able to eat some, but its program “understands” that it does not really perceive the amount of chocolate with any certainty; perhaps because it swallows all of it immediately without really taking a look. Intuitively, it does not then make a lot of sense to send a strong reward loss signal. Such a signal would be too much guesswork and would not provide a proper basis for learning better behaviour. In other words, since the system does not know for sure if there is frustration and how much, it should not send a strong frustration signal.<sup>19</sup> Thus, taking account of the uncertainty of the perception of reward would reduce suffering.<sup>20</sup>

### *The Buddhist concept of emptiness*

The perceptual kind of uncertainty has a central role in the later Mahayana school of Buddhism. While the “three characteristics” (impermanence, no-self, unsatisfactoriness) form the core of the Buddha’s original philosophy, later Buddhist philosophers found them somewhat simplistic. The emphasis shifted to the properties and limitations of perception and cognition, as opposed to characterizing the outer world. The inaccuracy of

---

was done in our frustration equation on page 158 by multiplying the perceived reward loss by the certainty of perception. See also footnote 10 in Chapter 13.

<sup>17</sup>Here we focused on the uncertainty in obtained reward. In an orthodox Bayesian interpretation, it may in fact not be possible to say that there is any uncertainty about the expected reward, since the expected reward is a subjective quantity, something purely defined by what the agent believes and expects. In contrast, in a “frequentist” interpretation, the expected reward is an objective quality in the outer world (how much the agent would get on average if it repeated the same action many times) so it can be misestimated, thus adding to the uncertainty of reward loss. Notwithstanding such theoretical arguments, I think it is clear that for biological organisms, understanding the real evolutionary value of, say, a piece of food may actually be a highly complex process involving a lot of learning and computation, so it can surely go wrong, as in the case of sugary food, which means there is uncertainty. I should also mention that the lowering of expectations discussed earlier in this chapter is different from the appreciation of uncertainty considered here since, in probabilistic terminology, mathematical expectation is completely different from variance.

<sup>18</sup>Let me just mention the great Chinese classic Huainanzi’s “The old man lost his horse”.

<sup>19</sup>The meaning of “weaker” signalling is a bit vague in this intuitive example. To make it more rigorous, we can consider, as an illustrative example, one of the simplest online learning tasks, namely linear regression. There, we minimize a quantity such as  $\sum_t (y_t - ax_t)^2 / \sigma_t^2$  where  $x$  is input,  $y$  is output,  $\sigma_t^2$  is noise level, and  $a$  is a parameter to be estimated. The magnitude of the error signal for each data point is proportional to the inverse of the noise level  $\sigma_t^2$ . Thus, for a high noise level (large uncertainty), the error signal is smaller. If the noise level is estimated separately for each data point (or time point  $t$ ), this will have the effect of reducing the error signal at time points where there is a lot of uncertainty as modelled by the noise level  $\sigma_t^2$ . The concrete algorithm used here might be what is called the delta rule; see Korenberg and Ghahramani (2002) as an example of a related if slightly more complex model. Mai et al. (2022) propose a closely related weighting in the context of reinforcement learning and RPE. See also footnote 15 in Chapter 5.

<sup>20</sup>One aspect of uncertainty which is not explicit in the frustration equation is the RPE due to changing predictions. Suppose one moment you think you will get a reward, but the next moment it looks like you will not get it. This decrease in expectation induces RPE and thus suffering (as explained in footnote 20 in Chapter 5). However, nothing may have actually happened, it was all just predictions in your mind. Importantly, such a change in prediction is only possible if there is uncertainty. If the future were certain, there would be no need to update your predictions, but because of uncertainty, the predictions change from one moment to another. Again, such suffering can be reduced if you realize that your predictions are uncertain; then the change in the predictions would be given less weight, as just explained in the text.

perceptions and beliefs became essential as part of the multifaceted concept of “emptiness”, widely used in Mahayana Buddhism—although rarely by the Buddha himself.<sup>21</sup>

Emptiness has many meanings. In the framework of this book, we can consider emptiness as an umbrella concept encompassing several of the ideas related to information-processing that we have seen in this book, in particular uncertainty, fuzziness, subjectivity, and contextuality. To summarize it in a single word well-known in Western thinking, we could call it “relativity”.<sup>22</sup> What the different aspects of emptiness have in common is that fully appreciating them should make us take the contents of our minds less seriously.

### *Giving up categories*

In particular, concepts and categories are considered fundamentally flawed in Mahayana philosophy. It proposes that the objects in the world do not really exist as separate entities, but are just part of a complex flux of perceptions happening in our consciousness. In this sense, there are really no separate objects or crisp categories in the world; they are purely constructions of the mind. Zen texts use the parable of confusing the moon and the finger that is pointing at the moon. Here, I would interpret this in the sense that the finger is a category, perhaps expressed by a word, that merely points at a phenomenon in the real world, that is, the moon. Ceasing to think in terms of categories and concepts, based on a recognition of emptiness, is something that generalizes the idea of reducing the certainty attributed to perception, or in fact, to your cognitive processes in general. It reduces frustration according to the logic given above for recognizing uncertainty of perception. Yet, there may be some processes more specific to categories.

If you admit that you’re not sure about what category some object belongs to, any further associations and generalizations have to be given up as well. For example, if I think that what somebody else just did was rude, perhaps I should not be so certain about such inference. To begin with, maybe my perception was incorrect: I may have completely misunderstood what he was doing, or what his goal was. From the viewpoint of contextuality, I might consider if in this particular situation, his behaviour was actually just right—or maybe I am in a foreign culture and don’t know the rules. From the viewpoint of subjectivity, I might wonder if other people

<sup>21</sup> Though see *Samyutta Nikaya* 22.95, where the Buddha clearly talks about a general emptiness of the Mahayana kind, while using a slightly different terminology: he does not use the word *suññatā/sūnyatā* which is the term usually translated as emptiness, and became prominent in later texts. See also *Majjhima Nikaya* 121 for a very different early view on emptiness, and Williams (2008b, p. 54).

<sup>22</sup> I am here referring to the common, non-technical definition of relativity, such as “the state of being dependent for existence on or determined in nature, value, or quality by relation to something else” (Merriam-Webster.com, accessed 24/1/2022). The interpretation of emptiness as relativity was initiated by Theodore Stcherbatsky, one of the earliest Western interpreters of Buddhist philosophy. Some commentators may prefer ontological interpretations of emptiness, but my treatment here sees it more as an epistemological quality, compatible with my computational approach. Emptiness actually has two different but related well-known definitions in Mahayana Buddhism (Williams, 2008b). First, there is the Yogācāra definition based on the “consciousness-only” thinking described in Chapter 12: All phenomena in the world are called empty because they are simply phenomena in the mind and constructed by the mind; in particular, any categories and concepts are merely mental constructs. This is rather similar to what we just discussed, except that in Yogācāra, such thinking can even be taken to a metaphysical level, denying the existence of the outside world—at least in some interpretations. Second, there is the Madhyamaka definition, where all phenomena are called empty in the sense that they are simply products of long causal chains, thus lacking any independent, intrinsic existence, and subject to change at any time. This is a very general definition that is ultimately supposed to contain most related properties described in this book or other Buddhist schools; it is surprisingly similar to the dictionary definition of relativity just given. For example, subjectivity of perception can be seen as a result of such causality because perception is causally influenced by the priors in the perceiver’s brain, and thus the percept does not exist independently (of the brain).



found his behaviour commendable and if it is just me who finds such behaviour rude. More generally, from the viewpoint of fuzziness, I might ask: How does one define rudeness anyway, is there a well-defined criterion?<sup>23</sup> In particular, any valence that you would typically associate with a category cannot be considered certain anymore. You may associate rudeness with a negative valence, but if you're not sure what is normal and what is rude, the negative valence cannot be generated anymore. What may be Epictetus's most famous quote says: "Men are disturbed, not by things, but by the principles and notions which they form concerning things."<sup>24</sup>

## Reducing self-needs

In the frustration equation above, we didn't have any terms explicitly related to self. Yet, self is obviously an extremely important concept from the viewpoint of suffering, as seen in Chapters 6 and 11. In our framework, self creates its own kind of frustration, by bringing aspects such as self-preservation, self-evaluation (or self-esteem), and control into play. As such, self-related suffering is covered by frustration equation as a special case. Many philosophical traditions such as Buddhism encourage reducing self-related thinking as a means to reduce suffering.

One case of self-related thinking concerns the self-evaluation system. In Chapter 6 it was proposed that a self-evaluation system constantly computes whether we have gained "enough" reward recently, looking at the relatively long-term performance of the system. (This long-term evaluation system is different from the one which computes the ordinary, short-term reward losses in the first place.) Such self-evaluation creates, as it were, another frustration signal on a higher level, in case the result of the self-evaluation is worse than some set standard.

Logically, there are three ways of reducing negative self-evaluations. The first is similar to the "conventional" approach we discussed above regarding ordinary frustration: it is to really gain a lot of reward, so that you surely reach the standard required. This is obviously easier said than done. Furthermore, such striving may not reduce suffering at all because gaining a lot of reward may increase the expectations in the future, resulting in insatiability on a "meta-level".<sup>25</sup> The second approach, in line with the main proposals in this chapter, is to lower the standard of expected reward. For example, the aforementioned philosophical viewpoint that everything is unsatisfactory should work here as well. If the system expects little reward even in the long run, the self-evaluation should not claim that the agent did not gain enough.

However, there is clearly a third option: shut down the system that evaluates your long-term success. Such

---

<sup>23</sup>Fuzziness is actually something whose effect on suffering we have not yet considered in detail, although it is an important concept—if not under this term—in relevant philosophical systems, such as Zen and the Greek (Pyrrhonian) Sceptics. Chapter 7 argued that while conceptual thinking uses crisp categories, many of the things in the world are fuzzy. While fuzziness is different from uncertainty, in our framework, nevertheless, it is related to the uncertainty of perception of reward loss. Valences of good and bad, and perceptions of rewards, are often based on categories of objects or events. (In fact, "good" and "bad" are categories themselves.) If you categorize a person's behaviour as "rude", you will perceive a negative reward, but maybe the behaviour was not actually that rude? If you categorize events which are only borderline rude as simply rude, that is a form of overgeneralization. You may be suffering unnecessarily due to your crisp-categorical thinking. The effect of fuzziness on suffering thus seems strongly analogous to the effect of uncertainty.

<sup>24</sup>*The Enchiridion*, Paragraph 5

<sup>25</sup>Actually, the theory in the previous chapters does not exactly lead to such meta-level insatiability. We saw in Chapter 5 how predictions are constantly updated, thus leading to insatiability. However, self-needs are not necessarily concerned with predictions but expectations of a different kind, as discussed in footnote 14 in Chapter 6. Still, it is possible that the expectations computed by the self-need systems are also updated based on past rewards, leading to meta-level insatiability.

a shut-down is possible by convincing yourself of the total futility of the self-evaluation. The Buddhist philosophy of no-self should be particularly useful here. Admitting the lack of control, even lack of free will, implies that there is little to evaluate: if the world gives us reward completely randomly, and we cannot influence it, what is the point in evaluating my actions and learning strategies? On a deeper philosophical level, if it is not me that actually decides my actions—but my neural networks, say—who is to be evaluated? Perhaps my neural networks and my body could still be evaluated, but not “me” really. On the other hand, what if “my” actions are ultimately determined by the input data, or the environment?

Suppose an agent were somehow able to shut down its self-evaluation system. It could be objected that such an agent with no self-evaluation would no longer be functional. However, even if the long-term self-evaluation were completely shut down, the system could still achieve most of its goals, and it will even be able to learn. Learning might just be slightly impeded because the learning system would not be optimally tuned to the environment. Thus, only “learning to learn”, a kind of meta-learning, would be shut down, while the agent would be perfectly functional otherwise, even without self-evaluation.

I should emphasize another crucial point about self-evaluation. As long as self-evaluation is based on evolutionary fitness, including what I called evolutionary *obsessions*, it does not actually make a lot of sense for us. It is too often based on criteria that are not in line with what humans should strive at according to various ethical considerations. We need better criteria to decide if our actions were “good enough”; criteria that would be more in line with what we consider a good human life should be about.<sup>26</sup>

Reducing the survival instinct, or behaviour and information-processing aiming at self-preservation, would seem to be equally useful for reducing suffering. Again, it could be objected that it is not good for the agent: such reduction may increase the probability of injury and even death. If I had no survival instinct, I might just happily go and pat a tiger I see in the jungle. This is a valid point, but we could still try to reduce the intensity of suffering incurred. In fact, religions and spiritual traditions invariably propose some method to cope with fear of death and mortality. Fear of death may often be unreasonable: I may even suffer from seeing a tiger on TV. Therefore, a moderate reduction in survival instinct might have mainly positive consequences. One method would be to reduce the mental simulations of injury and death; we will get back to this point in the next chapter where we look at reduction of simulation by meditation.

### *General reduction of self*

If we see self as the source of control, and then we admit uncontrollability as discussed above, this can be seen as a way of reducing the power of the self to influence our thinking. As far as self is about control, giving up control is, figuratively speaking, giving up part of the self. More precisely, it is rejecting part of the power that self-centered processing has on us.

Such “reduction of self” is a general principle that can take many forms. One approach is limiting the number of things belonging to self. For example, I could consider things that I think belong to me: perhaps my family, my house, my job and so on. If I think of them as “mine”, I invest them with a certain power because I

---

<sup>26</sup>As self-evaluation can use social comparison as a baseline (see footnote 2 in Chapter 6), it is also important to question the adequacy of such comparisons. For example, social media platforms may create unrealistically high standards regarding what one should look like and what lifestyle one should have, partly because such content is carefully selected and even fake. When adolescents compare their own life with social media, there may be a huge gap, which may create mental health problems as reported by Verduyn et al. (2015) (but see also Beyens et al. (2020)). How to reduce this kind of suffering: should people avoid using social media platforms? At the very least, it would be useful to understand the futility of such comparisons.

think I should be able to control them, as well as keep them intact. In other words, I think that they are in a sense part of myself; some would say I “identify” with them. Then, if anything bad happens to them, or anybody tries to take them away from me, I will have a strong negative emotion as if my self were threatened—and in a sense the intactness of my person or self *is* threatened.

It is clear how one can reduce suffering coming from such possessions: as a first approach, just own fewer things. If you have very few things which you consider yours, it is less likely that you will experience them breaking down, being stolen, or getting lost. Many spiritual traditions do recommend giving up most of your material possessions. Further, you can try to change your attitude towards such external parts of yourself. Epictetus proposes that you should think of all your possessions, your family etc., as not really belonging to you, but as things that have been given or lent to you:<sup>27</sup>

Never say of anything, “I have lost it”; but, “I have returned it.” Is your child dead? It is returned. Is your wife dead? She is returned. Is your estate taken away? Well, and is not that likewise returned?

Finally, the reduction of self can be approached from the viewpoint of reducing thinking in terms of categories. Typically, I divide the world into things that are part of myself, and things that are not part of myself. This is how I construct the category “self”. Like with other categories, it would be useful not to take this category too seriously, and understand its fuzziness and arbitrariness, or emptiness. “Self” can be seen as the ultimate category that should be deconstructed and given up. Such giving up of the whole category of self, in a sense, encompasses all the other aspects of no-self philosophy described above. If the very category of self does not exist, or, to put it simply, if self does not exist, what would be the point in self-preservation or self-evaluation, or any attempt to control? Any such self-related thinking should vanish if the underlying category of “self” is given up. The Buddha said that when a monk is advanced enough, “any thoughts of ‘me’ or ‘mine’ or ‘I am’ do not occur to him”.<sup>28</sup> This is the most general way of reducing suffering based on no-self philosophy.

## Reducing desire and aversion

While so far we have focused on reducing the frustration of desires, many philosophical traditions propose that desires themselves should be reduced—as always, this includes aversions. In Buddhist philosophy of the Theravadan school, it is traditionally the main focus of the training, and it is the main point of the Buddha’s teaching as expressed in the Four Noble Truths. After describing what suffering is (quoted on page 19), he proposed that it is born of desire, and that one can be liberated from suffering by eradicating desire following a path of meditative and other practices.<sup>29</sup> Epictetus was equally clear about the importance of not having desires or aversions, especially towards things we cannot control:<sup>30</sup>

<sup>27</sup> *The Enchiridion*, Paragraph 11. See also *Samyutta Nikaya* 22.33, where the Buddha takes this approach to the extreme in the sense that he recommends abandoning everything, including any aspects of your mind.

<sup>28</sup> *Samyutta Nikaya* 35.205

<sup>29</sup> For completeness, I will briefly describe the Buddha’s Four Noble Truths in their entirety. They can be seen as a psychological theory of why suffering comes about and how it can be avoided. The four truths are a logical sequence: 1) All phenomena (i.e. external objects, perceptions, feelings, thoughts, etc.) in the world are unsatisfactory in the sense that they have the potential to produce suffering. 2) Suffering is produced by desire for any of these phenomena (or desire to avoid any of them, i.e., aversion). 3) Suffering disappears if desire is eradicated. 4) Desire can be eradicated by following a certain combination of meditation techniques, philosophical attitudes, and ethical behaviour. (For references, see footnote 20 in Chapter 2.)

<sup>30</sup> *The Enchiridion*, Paragraph 2; see also *The Discourses* I:4 by the same author. The same point was further made by Epicurus—who seems to have been seriously misunderstood. Epicurus proposed that there are a few desires which need to be satisfied since they

Remove aversion, then, from all things that are not in our control (...) But, for the present, totally suppress desire: for, if you desire any of the things which are not in your own control, you must necessarily be disappointed; and of those which are, and which it would be laudable to desire, nothing is yet in your possession.

Humans can indeed reduce frustration simply by giving up some unnecessary goals: you don't really need a fancy car. It is possible to consciously decide not to strive for certain goals, and we can modify our desires to some extent without any special techniques. In our framework, this in particular means reducing intentions, i.e., commitment to plans, also called attachments in Buddhist terminology. (Intentions can, in fact, be easier to reduce than desires themselves, as may be intuitively clear and will be discussed in more detail in the next chapter.) Suffering will then be reduced since each goal could potentially lead to frustration. If there are no desires and no goals that need to be achieved, frustration will not appear and nor will suffering. As such, reduction of desires is a central mechanism through which reduction of frustration is possible.

Many ideas in this chapter can be seen as mental techniques serving the very goal of reducing desires. Consider, for example, reducing expected rewards as considered above: why would the agent want anything if it has arrived at the conclusion that the expected rewards are zero, or very small? Likewise, desires will be reduced by adopting the belief that many desires are pointless and even bad for you, they are just evolutionary obsessions. As such, reducing desires is closely related to the earlier ideas of facing uncertainty, uncontrollability and unsatisfactoriness, and in fact, in a traditional Buddhist account, the main justification for such philosophical attitudes is precisely that they reduce desires.<sup>31</sup>

There are also special techniques to reduce desires. One example is choosing to pay attention to good things that one already has, instead of things that one might obtain. This reduces desires and the tendency of insatiability; it is central in mental exercises based on gratitude.<sup>32</sup> Further, Epictetus proposed a rather extreme form of contemplation of impermanence, namely contemplation of death.<sup>33</sup>

---

are both natural and necessary: Food, water, and shelter; these desires are also easy to satisfy. In contrast, desire for money, power, fame etc. are unnatural and unnecessary; they are also insatiable. Optimal “pleasure” is obtained by rejecting desires which are not natural and necessary. See Epicurus's *Letter to Menoeceus*, Hadot (2002, p.34) and (Konstan, 2018) on the insatiability/satiability (or satisfiability) distinction.

<sup>31</sup>Let me try to make the links to the other ideas in this chapter explicit. First, reduction of certainty attributed to perception reduces desires since you don't actually know for sure whether the object of your desire really gives reward—or is even there. Second, if you cannot control anything, what would be the point in wanting, let alone planning, since rewards and goals cannot be attained? Seeing the insatiability of the desires should also lead to the conclusion that their total fulfillment is impossible in the long run, so the desires should be dropped as futile; seeing desires as evolutionary obsessions means realising they can even be bad for you. Self-needs, in particular self-evaluation, can be considered as forms of desires in this context, and the same ideas apply to them.

<sup>32</sup>(Emmons and Shelton, 2002)

<sup>33</sup>Quote from *The Enchiridion*, Paragraph 21. It is easy to see why contemplation of death would reduce desires, and in particular planning and intentions. Presumably, contemplation of death reminds you that you just might die tomorrow, even if that is not very likely. By some kind of availability heuristic (see footnote 12 in this chapter), that reminder will increase your estimate of the probability of dying soon, which implies that you don't have much time to obtain rewards left. So, their expected value is low, and any planning is less useful and highly restricted by this time horizon. (Such a reminder might also lead to gratitude for being alive; understanding the mechanisms of gratitude is, however, outside of the scope of this book.) Buddhist practices also include contemplation of death; it may serve slightly different purposes (Anālayo, 2003, p. 155), but the classical manual *Visuddhimagga* (Chapter VIII, 41) links it directly to “disenchantment” and “conquer[ing] attachment”. Nevertheless, some psychological research based on the Terror Management Theory claims that reminding people of their mortality may, in fact, increase their willingness to consume (Kasser and Sheldon, 2000); see also Burke et al. (2010); Gao et al. (2020). It remains to be understood why such quite opposite effects can be observed.

Let death and exile, and all other things which appear terrible be daily before your eyes, but chiefly death, and you will never entertain any abject thought, nor too eagerly covet anything.

Yet, there are also desires that are really “hot”, hard-wired, and difficult to modify, let alone reduce, based on the rather purely philosophical or intellectual considerations presented in this chapter. What is needed are special techniques to work on deeper levels of the mind than philosophical thinking. Meditation is one such method, as we will see in the next chapter.<sup>34</sup>

## How far should reducing desires and expectations go?

An objection could be raised at this point: the thinking discussed in this chapter seems depressing. One may ask whether it does not lead to complete inactivity, and, indeed, to some kind of depression. A diagnostic criterion of depression is “markedly diminished interest (...) in all, or almost all, activities”,<sup>35</sup> which sounds a bit like having substantially reduced reward expectations, and having few desires. The fundamental question is: can such reduction of desires and expectations go too far? It is related to the question of how far should such lowering of expectations go. Is it enough to admit the *actual* levels of uncertainty and uncontrollability, or should we go further and consider things even more uncertain and uncontrollable than they really are, thus lowering expectations even more?

If our only goal were to reduce suffering in the agent, we could simply program it to assume that everything is completely uncertain and completely uncontrollable. Then, the agent would expect zero reward, or very little, in any state or from any action. It would have virtually no desires, which is a consequence of such a radical reduction of reward expectations. Is this a good way of programming an agent?<sup>36</sup>

Claiming that Buddhist training can lead to something akin to depression is, in fact, a well-known point of criticism, and similar arguments have actually been raised against Buddhism throughout its history. I think such criticism is not very relevant because it considers an extreme case, which is unlikely to be achieved by most people practising such systems. Perhaps the point is that most people living in a modern industrialized society simply have too many desires, and it would be better for them to have fewer of them. This would explain why people engaged in Buddhist training tend to get happier when they reduce reward expectations. It may be

<sup>34</sup>In addition to reducing desires, Stoics proposed another approach to working with desires. It is also possible to align one's desires with what can be more easily achieved in the world, instead of eliminating them. Epictetus takes this idea to the extreme by suggesting: “Don't demand that things happen as you wish, but wish that they happen as they do happen, and you will go on well.” (*The Enchiridion*, Paragraph 8) If the only thing you want is that things happen as they do happen, how could there be any frustration? Your wishes will always be fulfilled.

<sup>35</sup>DSM-5 diagnostic criteria

<sup>36</sup>The traditional Buddhist viewpoint tends to emphasize that people are mistaken about the level of control and permanence, and it is enough to correct their “ignorance” or “illusions”. On the other hand, consider a super-intelligent agent which has no constraints regarding data or computation. It would presumably estimate uncontrollability and uncertainty correctly and accurately, without any illusions. But it would still have reward losses, and those reward losses might not even be particularly small, especially if the outside world is difficult to control (perhaps due to strong physical constraints in the ability of the agent to manipulate it) and exhibits a lot of randomness. So, it is not clear if suffering would be very much reduced by correcting “illusions” in the sense that the agent learns to make “optimal” inference (in the sense of probabilistic AI theory) with infinite data and computation. I would assume that the real goal of such Buddhist practice may rather amount to adopting reward expectations which are lower than what is objectively true. In this case, it would lead to increased happiness at the expense of slightly suboptimal inference—but note that such “suboptimality” refers only to the lack of optimality in maximizing rewards.

irrelevant to ask what might happen in the extreme case where they completely annihilate all their desires—which is a feat even most meditation masters are incapable of. Buddhist philosophy actually emphasizes the “middle way”, or moderation, which sounds like a good idea here as well. The situation might be different for Buddhist monks or nuns engaged in full-time practice for many years; but they are following a very special lifestyle, which is specifically designed to be compatible with having very few desires.<sup>37</sup>

On the other hand, there is certainly something fundamentally different between a depressive state and a mental state where the unsatisfactoriness of the world is seen from a Buddhist perspective. If an agent concludes that none of its desires are going to be fulfilled and it will never receive any reward, that gives in itself no reason for a negative feeling or valence. The agent would just rationally decide that no desires are worth pursuing, it would not engage in goal-oriented action, it would predict zero rewards in the future, and, in fact, it would suffer *less* since there is no frustration. If humans tend to get a negative feeling after seeing that the world is fundamentally unsatisfactory, it must be because there is another “higher-order” desire, presumably coming from the self-evaluation system treated in Chapter 6.

A depressed person, in particular, finds the very unsatisfactoriness of the world frustrating, and *wants* to find satisfaction or reward in various kinds of seemingly pleasurable objects and activities. In our framework, we would say that she is frustrated in terms of her self-evaluation, as she sees that she gets less reward in the long run than she “should” according to some internal standard. The self-evaluation system may indeed conclude—based on a superficial calculation—that if it is true that no goals will be reached and no reward will be obtained, there must be something wrong with the agent, and a negative meta-learning signal should be generated, which would be felt as suffering. However, I think an important point in the Buddhist philosophy of unsatisfactoriness is that if the self-evaluation system sends such a negative signal, it is simply malfunctioning. Clearly, the realization of the total unsatisfactoriness of everything should also influence the self-evaluation system. The self-evaluation system should set its expectations, or standards of “acceptable” reward level, very low, even zero. The self-evaluation system cannot rationally claim that the agent is not getting enough rewards if it believes itself that no rewards can possibly be obtained! As such, Buddhist philosophy proposes that there is no need to be frustrated about any long-term lack of reward, nor is there any need to make any negative self-evaluation; not getting much reward and not reaching one’s goals is natural and unavoidable.<sup>38</sup>

### *Is frustration not needed for learning?*

Another objection that could be raised against the philosophy presented here is that it may not be useful to reduce frustration since the frustration signal is useful for learning. Human beings seem to be trapped in a situation where they need frustration to learn, while they suffer from it. That may sound like a dilemma with no satisfactory solution. However, I’m not sure there is any real dilemma here. One reason is that, as discussed in Chapter 5, many of the rewards we are programmed to receive are actually rather useless “evolutionary

<sup>37</sup>A famous counterexample to my optimism happened during the Buddha’s life, when several of his disciples committed suicide after intensively engaging in a particular exercise: reducing carnal desires by contemplating the loathsomeness of the human body (*Samyutta Nikaya* 54.9). The Buddha realized his mistake and changed his teaching accordingly. That was a case of using loathing as a meditation technique to reduce reward expectations and desires, which is extremely rare in current Western meditation practice.

<sup>38</sup>We will see, in the next chapters, two more related points which relativize the importance of reducing desires by reducing reward expectations. First, moving to the level of meta-cognition may mean that reducing desires is less important than seeing them from a meta-cognitive level. Second, in the case of the bodhisattva ideal of the Mahayana schools, the desire for reducing other people’s suffering may actually be good.

obsessions”; frustrating them may not teach us anything useful, if it is not the very futility of those rewards. The same is true from the viewpoint of insatiability: why should one try to learn how to better satisfy desires which cannot be satiated anyway?

Furthermore, Chapter 12 proposed that a large part of the problem is actually how frustration is made conscious even it didn’t need to be; learning from frustration could, in principle, happen on an unconscious level. It may seem that not much can be done about this, but in fact, an intervention is possible, as will be seen when we talk about meta-awareness in Chapter 15.

Another point is that while understanding the uncertainty of all perceptions reduces frustration, it may actually *improve* learning and make us more “intelligent”. Uncertainty and uncontrollability are real properties of the world, but we may have been grossly underestimating them because our basic cognitive apparatus is not good at handling them. Thus, learning to better appreciate the uncertainty and the uncontrollability is a useful learning process even from the viewpoint of trying to optimize rewards in the world. Based on these arguments, I think that while it may be meaningful to claim that not all frustration should be removed, *most* of it can still be removed without making learning any worse.<sup>39</sup>

### *Contentment and freedom*

To conclude this chapter, let us look at the absence of desires from more positive viewpoints. First, the absence of desires can be expressed as *contentment*, in the literal meaning of being content with what one has and not wanting more. Insatiability of the desires, in particular, implies that the agent cannot be content: it always has to search for more rewards. That leads to endless frustration, which is why reducing desires is essential. In a stronger form, contentment may further turn into a feeling of gratitude.<sup>40</sup>

Another well-known positive interpretation of having no desires is *freedom*. The condition of a human being has been described as being a “puppet of the gods” by Plato,<sup>41</sup> meaning that “affections in us are like cords and strings, which pull us different and opposite ways”. We have to remove those cords and strings if we want to be free, instead of being enslaved by the interrupts and unconscious action tendencies. Epictetus summarizes the idea in a way that is, yet again, in complete harmony with the Buddha’s philosophy.<sup>42</sup>

Freedom is acquired not by the full possession of the things which are desired, but by removing the desire.

---

<sup>39</sup>In fact, if somebody argues that frustration is actually good since it enables learning, the question arises as to why frustration is then painful. If learning is “good” and should be encouraged, and frustration is needed for learning, frustration should logically feel pleasant, not painful. The fact that frustration is painful means that at least in some sense and to some extent, it has been deemed to be bad for you. This may only hold from an evolutionary viewpoint, and only under some “normal” conditions, though. As a thought experiment, suppose frustration felt good, perhaps because you have become so thoroughly convinced about the utility of the ensuing learning that you are able to override millions of years of evolution. Then, you would presumably try to fail in everything you do—it feels good and you will consider that good feeling as a reward. You might learn a lot from such failures, although if you fail without even trying hard, the utility for learning might be meager. In any case, you would not get any reward; you might starve, die young, and would not produce any offspring. This (admittedly not very rigorous) argumentation suggests that frustration has to be evolutionarily made painful, at least to a certain extent. However, this argumentation was based on an extreme case. Perhaps there is an optimal amount of frustration which is not zero; perhaps it would be possible to detect circumstances under which frustration is good while it is usually bad. I leave this for future research.

<sup>40</sup>(Emmons and Shelton, 2002)

<sup>41</sup>Plato, *Laws*, Book I. These puppets are different from those that Plato talked about in his more famous cave allegory. See also Marcus Aurelius’s *Meditations* II.2.

<sup>42</sup>*The Discourses* by Epictetus, IV:1. Peacock (2018) formulates the goal of early Buddhist philosophy as freedom from “reactive patterns” triggered by valences.

## Chapter 15

# Retraining neural networks by meditation

The preceding chapter presented several directions in which information-processing should be changed in order that suffering is reduced. We also saw some practical suggestions for reprogramming, such as seeing the uncertainty and uncontrollability of the world and reducing desires and self-needs. This will eventually lead to a reduction in reward loss, frustration, and suffering. Yet, the account of the preceding chapter may be rather unsatisfactory for some readers: It seems to be asking the impossible, at least in the case of mere humans. The goal is to change some fundamental beliefs about the world and your mind. How is one supposed to become so thoroughly convinced about, say, the uncontrollability of the world that one is not disturbed by the loss of, say, one's job or house? Is it not simply "human" to think otherwise? How can you actually reduce expectations of rewards, belief in the certainty of perceptions, and so on?

Crucially, what we need are changes in neural associations which work on an *unconscious* level, and that is notoriously difficult. We need to develop practical methods for retraining the neural networks in the human brain.

In this chapter, I consider meditation, or mindfulness training, as a method that can radically boost retraining of neural networks, compared to straightforward attempts to change thinking at the conscious level. It also turns out that meditation has further benefits, such as reducing "hot" desires, reducing simulation and developing metacognition. I will not go too much into the practical details of any such training methods, on which hundreds of manuals have already been written. Rather, I discuss general principles on how they work, largely interpreting them in the information-processing framework of this book.

First, I discuss how the meditation methods can be seen to speed up learning from new input, thus enhancing the methods of the preceding chapter. Second, meditation can be seen as reducing two terms in the frustration equation on page 158 that we did not yet consider: the amount of attention paid to reward loss and the number of times the reward loss is perceived or simulated; these are also related to the top-row green boxes in the flowchart in Figure 13.1. In fact, emptying the mind by meditation clearly reduces simulation, and meditation almost inevitably seems to develop a metacognitive attitude which changes the attention paid to reward loss. A third benefit of meditation is that it enables stopping the processing chain in the flowchart in Figure 13.2 by increasing conscious control over interrupting desires. Finally, I discuss how meditation can be interpreted as mental relaxation.



## Contemplation as active replay

The fundamental problem with the approach of the preceding chapter is that a conscious decision to think in a different way often has little effect on what unconscious neural networks do. A conscious decision may not even really change future conscious thinking since it may be overridden by the unconscious networks. That is why reprogramming of the brain must include some kind of retraining the unconscious neural networks.

From a dual-process perspective, the problem to be solved here is how the conscious-symbolic-explicit system can change a mental association which is actually encoded in both the two systems. For example, it might try to create an association between “I” and “impermanent”, being inspired by classical Buddhist philosophy. However, what really matters is changes in the unconscious-neural-implicit system, because it is that system that computes values, expected rewards, and reward losses. So, how can the explicit system force a change in the implicit one? Transfer of knowledge or learning between the two systems is difficult. While one may have a clear understanding that everything is impermanent on a conscious level, it is not easy to transfer this understanding to the unconscious neural networks.

As a first approach, we could use techniques that I here call *contemplation*. That means a constant conscious repetition of selected thoughts. For example, it can be contemplation of the characteristics of impermanence, uncontrollability, and unsatisfactoriness, possibly combined with some object—such as “I” or my “self”—whose impermanence or other property one wants to learn. The constant repetition of such thoughts on a conscious level should slowly modify the unconscious associations used to compute the perceptions and replay. Some kind of Hebbian learning is likely to construct an association, even on the basic neural level, between the different concepts, such as “I” and “impermanent”. Reading books on Buddhist or Stoic philosophy, as well as later thinking about their contents, can also be seen to be such contemplation.

The mechanism working here is what I would call *active replay*: The explicit system uses the mechanism of experience replay (see Chapter 9) to make the implicit system learn whatever the explicit system wants. That is, the explicit system in your brain can select thoughts in the form of linguistic sentences or visual images of events—possibly imaginary—and replay them. It can do that repeatedly, thus replaying selected items many times. Such replay will change your neural networks—that is in fact the very point in replay. What is special here is that the explicit system chooses what to replay, thus “teaching” the neural networks, while in ordinary replay, the material would be selected by the implicit system itself.

Such training may seem rather different from modern meditation instructions, but it seems to have been an essential form of practice in the Buddha’s times, and emphasized by some modern Buddhist meditation teachers as well.<sup>1</sup> When the Buddha was asked for meditation instruction by monks entering a solitary retreat, he would often tell them to contemplate on impermanence, no-self, or unsatisfactoriness, sometimes linking them all together in various causal chains. For example, he would advise:

You should abandon desire for whatever is impermanent. And what is impermanent? The eye [and visible forms etc.] is impermanent; you should abandon desire for it.<sup>2</sup>

Forms [i.e. anything that is seen] are impermanent. What is impermanent is suffering. What is suffering [i.e. unsatisfactory] is nonself. What is nonself should be seen as it really is with correct

<sup>1</sup> See e.g. Anālayo (2003, p.103-104); Mahasi (1996)

<sup>2</sup> *Samyutta Nikaya* 35.76; see also *Samyutta Nikaya* 35.32; *Samyutta Nikaya* 35.162

wisdom thus: “This is not mine, this I am not, this is not my self.”<sup>3</sup>

We don’t know much about the details of how such contemplation was practised in the Buddha’s times. Presumably, it was combined with meditation to speed up the learning—more on that below.

The fundamental problem with such basic contemplation is, in fact, that the learning process can be very slow and inefficient. One reason is that it has the same characteristics as learning in neural networks in AI. As we saw earlier, neural network learning requires a large number of repetitions of input, which changes the neural connections little by little, using some mechanism related to stochastic gradient descent or Hebbian learning. So, retraining neural networks by contemplation requires a huge number of repetitions.

Moreover, transferring learning from the explicit to the implicit system is hampered by the fact that the representations and computations in the two systems can be quite different, as we have already discussed. Suppose that your explicit system repeats the word “impermanence”, in an attempt to contemplate on that property. How are your primitive, lizard-level neural networks supposed to understand what that means? Such neural networks do not operate with words or abstractions but on representations related to sensory input. There is a kind of a communication barrier between the two systems, and contemplation may not be able to cross it very well.

The situation can be somewhat improved if the explicit system imagines events or episodes and replays them as real sensorial input such as images, instead of merely in verbal and abstract form. When you read a story or a simile in Buddhist literature and vividly imagine it happening, that does provide more natural input to your neural networks. Or, the explicit system can imagine future events from the viewpoint that an action plan is likely to produce frustration, as in Epictetus’s Roman bath example (see page 164).<sup>4</sup>

## Mindfulness meditation as training from a new data set

A crucial improvement to such contemplation practices is what is called meditation in the modern context. Mindfulness meditation in particular is a technique that can influence neural networks more efficiently than simple contemplation. Mindfulness meditation can incorporate many of the goals described above, such as realizing uncertainty and uncontrollability.<sup>5</sup>

Typical instructions of mindfulness meditation emphasize objective observation of any contents that appear in your mind, that is, mental phenomena. In particular, that encompasses anything that your senses perceive, including the “internal sense” of thinking and imagination. If you hear something, you acknowledge hearing it, if there is a bodily feeling in any part of your body, you recognize that you have a bodily feeling, and so on. Such observation is done, as far as possible, passively without interfering with the sensory process or the physical source of the perceptions (for example, without moving your body to change bodily feelings).

<sup>3</sup> *Samyutta Nikaya* 35.4. See also Williams (2008b, p.79) for a description of similar practices in a Mahayana context.

<sup>4</sup> One way of improving learning would be to adapt the contents of the contemplation to each individual based on their personality and temperament. While this is rarely done in a Buddhist context, the classical Buddhist meditation manual *Visuddhimagga*, for example, does include such instructions (Chapter III, 74).

<sup>5</sup> For introductory books to mindfulness meditation, see e.g. Gunaratana (2010); Kabat-Zinn (2012); for an attempt at a definition, see Bishop et al. (2004). In this book, the term mindfulness always refers to mindfulness meditation, often seen as a training or a learning process (instead of a state of mind, or a long-term psychological trait, which are alternative uses of the term). In terms of a well-known typology of meditation practices (Lutz et al., 2008; Dahl et al., 2015), what is emphasized here is the “open monitoring” aspect of the practice. Meditation is thus virtually synonymous with such terms as insight meditation or vipassana as far as this book is concerned.

The contents should be observed from an external perspective, as if from a distance, and without judging the contents to be either good or bad.

There are a number of techniques to make such observation easier by regulating the attention of the meditator. Basic meditation instructions typically start by recommending sitting in a comfortable posture and then provide one particular technique for attention regulation. A classical one is focusing on observing the breath, possibly reinforced by mentally counting the breaths. (Alternatively, the focus might be a visual target, or a particular word or phrase that is mentally repeated.) Such observation of breathing should be seen as simply a technique whose goal is to enable better observation of the mental phenomena, and indeed simply counting breathing may sound like an absurd exercise if the actual purpose is not understood. The purpose is to facilitate observation of the mind by making it relatively empty; observing mental phenomena is very difficult if the mind is full of different kinds of thoughts and perceptions. Furthermore, emptying the mind has several direct benefits as well, in particular reduction of simulation as discussed below.

The exact mechanisms of mindfulness meditation are far from being understood, but some of them can be understood by the framework presented in this book, as we will see next.<sup>6</sup>

### *Direct input to train neural networks*

The most crucial mechanism at play may be that the meditator learns largely the same things as in the contemplations above but in a more efficient way. I suggest the reason why meditation is more efficient than what I called active replay above is that there is no longer any need to transfer information between the two systems (conscious thinking and neural networks, roughly speaking). Instead, the practitioner observes characteristics such as impermanence first-hand, in real sensory input or imagined sensory content. Then, neural network learning can proceed in a completely natural manner, largely bypassing linguistic constructs and conceptual thinking.

In other words, during meditation, the sensory systems *directly perceive* how things are, say, impermanent by observing how those things change and disappear. Likewise, the control systems, based on sensory input, find by themselves that attempts to control fail when wandering thoughts invade the mind, for example. Thus, the neural networks learn directly from such natural input. This is in stark contrast to contemplation, where the difficult part is to transform concepts and words into something that can train neural networks, and replay does this in a somewhat contrived way. Neural networks learn best from real sensory input, so it is crucial here to enable them to do exactly that. Such observation is eventually extended to all the aspects discussed in the preceding chapter.

The key trick here is to select the right data to input into the neural networks. As discussed in Chapter 10, selection of input data is an essential part of the perceptual system, in terms of the multi-faceted phenomenon called attention. That is why regulation of attention is a central part of any meditation method: in mindfulness meditation, you typically start by focusing your attention on observing your breath. It is in fact possible to get useful input data from the breath itself, if you do it with a special kind of attentional focus. While the practice may start by simply observing the breath in a general manner, eventually, you can start observing its specific

---

<sup>6</sup>For previous proposals and reviews on the mechanisms of mindfulness, see e.g. (Hölzel et al., 2011; Grabovac et al., 2011; Williams, 2008a; Shapiro et al., 2006; Teasdale and Chaskalson, 2011b; Vago and David, 2012; Baer, 2003; Garland et al., 2014; Wielgosz et al., 2019). My account here is particularly computational, and emphasizes the learning of new attitudes which could be called philosophical, but crucially they are to be adopted in the neural networks as well. The emphasis is on understating uncontrollability, uncertainty, and unsatisfactoriness.

aspects in light of the theory of the preceding chapter. For example, you observe the impermanence of breath, how it is changing all the time from an in-breath to an out-breath—this is a classical Buddhist exercise. That means your attentional system selects your sensory input to consist of observations of your breath, and more precisely, any aspects of your breath related to permanence or lack of it. This is how your neural networks get a lot of good data pertaining to that particular property, and they learn to perceive the impermanence much better than they would by any kind of abstract contemplation based on linguistic concepts.<sup>7</sup>

So, the explicit system in a sense “teaches” the implicit system, and the teaching happens by means of the attentional system. Direction of attention is, to some extent, under conscious control. So, the symbolic or thinking part of the brain can just tell where the implicit system should be looking—this is only partly a figure of speech— and it does not need to really input anything to the implicit system, unlike in the case with replay. It is a bit like a professor telling students to read a book; she does not then need to give a lecture herself.

### *Realizing how the mind wanders*

Another important example of such direct input is observing how often and easily the mind starts wandering. As we have seen in Chapter 9, sustained attention is difficult, and after a while, the mind often starts wandering, and various daydreams fill the mind. Frequent occurrence of such mind-wandering is extremely salient to anybody who tries to focus on breathing or a similar meditation object. Realizing how difficult it is to focus on breathing gives a direct view into how uncontrollable the mind is. If you systematically observe how automatically your mind starts wandering, you will gradually be convinced—and so will your neural networks—that you cannot control even your own thinking, at least not completely. After all, wandering thoughts are, by definition, a failure of controlling your mind.

Such observation may also convince you that there is no self, no central executive, and perhaps then no free will. Under ordinary circumstances, if I decide to plan what I will do tomorrow, I may have a clear feeling that it is “me” who is doing the planning. However, after observing how planning happens automatically in wandering thoughts, I may be forced to admit that the plans are something that “I” did not create. You may even start having doubts about the correctness and certainty of your thoughts, since they seem to be something that just appears in the mind, and you have little idea why they appear or where they come from. Thus, uncertainty about your thoughts can be taught to the neural networks as well.

To interpret this learning process in the framework of the frustration equation, what happens is that the *unconscious neural networks themselves*—and not just the conscious and/or symbolic thinking systems—will *learn to reduce the expectations* of any rewards. This happens through your neural networks learning that the world is uncontrollable and uncertain, which necessarily reduces their expectation of future rewards according to the logic of the preceding chapter.

Many further meditation techniques can be seen as such attentional selection of particular direct input. One classical Buddhist technique is to focus on the ending of any pleasurable feeling. This enables seeing first-hand how pleasure, and in general any effects of rewards, are fleeting and thus worth less than might be expected. Thus, you will learn the impermanence and unsatisfactoriness of all mental phenomena in a particularly efficient way.

---

<sup>7</sup>Such selection of input can be further improved by controlling one’s media consumption as well as by choosing a suitable lifestyle and social environment. Buddhist monastic training provides a rather extreme example of such choices.

*Extinction of aversive responses*

In a slight variant of the logic above, mindfulness meditation can also help directly change associations related to specific emotions. An important example is fear extinction. Extinction is the opposite of classical conditioning: It means that when the predictive stimulus (e.g. the bell for Pavlov's dog) is presented *without* the other stimulus (the food for the dog), the conditioning weakens. If the bell is presented without the food many times, the dog learns that the bell does not predict the food anymore, and the conditioning is eventually extinguished.

Suppose you have learned to associate a fear reaction with your boss by classical conditioning. Perhaps that was based on a single episode, and the association is not valid anymore, so it would be nice to be able to let such a fear reaction be extinguished. Unfortunately, extinction is often very slow—just like any neural network learning—but this can be improved by mindfulness training. The trick here is that you create completely new data, going beyond simply selecting input from existing data as above, but still feed it directly into your neural networks.

It turns out that mindful meditation tends to make people relaxed and feel good (possible reasons for this will be discussed later in this chapter). So, if you recall the unpleasant episode with your boss many times, but always stay in such a pleasant, calm meditative state, extinction is more likely to happen. Thoughts about unpleasant situations will be increasingly associated with a general feeling of calm; the image of your boss will be associated with relaxation and feeling good in the whole body. This will help override the fear association.<sup>8</sup>

**Speeding up the training**

Unfortunately, such meditation training is still rather slow, even if it improves on simple contemplation. In fact, slowness of training is a ubiquitous problem with neural networks, as already pointed out in Chapter 4. Even though with mindfulness meditation, we have a new source of more direct and natural data for learning, the neural networks still need large amounts of input data, and a lot of meditation practice is needed. Fortunately, the amount of training and effort required can be further reduced by further techniques.

*Increasing the plasticity of the brain*

One central principle here is increasing the *plasticity* in the brain. Plasticity is the biological term for the capacity of neural connections to change and thus to learn. Plasticity in the brain's neural networks is by no means granted, nor is it a constant quantity. If by some suitable tricks, such learning capacity could be increased, the learning process would take less time. A large amount of neuroscience research has been dedicated to finding different ways to increase plasticity.

Sensory deprivation seems to be one useful trick; it has indeed been shown to increase plasticity, at least in rats and cats. It may be rather common sense that if your brain has had little stimulation for a while, it will better concentrate on any new task. It turns out that its learning capacities are also increased.<sup>9</sup> Mindfulness meditation in itself can be seen as imposing sensory deprivation, since it is usually conducted in a quiet

---

<sup>8</sup>On the general idea of extinction by mindfulness and exposure, see (Baer, 2003; Hölzel et al., 2011); on relaxation and positive affect, (Carmody, 2015). An interesting question here is whether extinction needs attention and/or awareness, and would thus be greatly facilitated by mindfulness; the results are not very clear-cut on this point (Kwapis et al., 2015; Han et al., 2003; Weike et al., 2007).

<sup>9</sup>(He et al., 2006; Duffy and Mitchell, 2013)

environment with eyes closed, or at least there is nothing much happening in the visual field. In some meditation schools, an even stronger form of sensory deprivation may be imposed in the form of silent retreats. Such retreats often entail minimization of any kind of sensory stimulation: the participants don't go out of a prescribed enclosure, they don't watch TV or use the internet, and obviously they don't talk to each other. In several discourses, the Buddha recommended such deprivation, together with meditative concentration, because it makes the mind "pliant" and "malleable". Then, the meditator is better able to gain insight into, for example, the uncontrollability and uncertainty of existence, as well as better able to learn from those insights.<sup>10</sup>

Plasticity can further be increased by restriction of food intake, which is another typical characteristic of ascetic training in spiritual traditions. Paradoxically, it can also be increased by the very opposite of sensory deprivation: enriching the environment. In animal experiments, that might mean allowing the animals to live as groups in large, spatially complex cages which are equipped with toys and running wheels. In humans, similar results are obtained by aerobic exercise, as well as action video game playing. Whether such methods could be used to improve meditation practice is a very interesting question for future research.<sup>11</sup>

Plasticity can also be increased by drugs, such as the antidepressant fluoxetine (aka Prozac). A large amount of research is currently being conducted on new drugs that would increase plasticity even more, and with minimal side effects. The huge impact such drugs could have on society is obvious.<sup>12</sup>

In fact, you may be wondering why plasticity is such a bottleneck: Why hasn't evolution made our neural networks learn much faster? The reason seems to be that some limitation of plasticity in the brain is useful to prevent new information from overwriting old information too easily.<sup>13</sup> So, it may not be wise to increase plasticity too much, because it could lead to too much forgetting of previously learned information. This is hardly a problem with meditation-based interventions, but with drugs, such negative side-effects might be real.

In principle, an AI has much more freedom in how it changes the results of its learning, and the amount of "plasticity" could be made infinite by design. Thus, an AI could get rid of a bad habit or a harmful association in a split-second, by just removing or changing some connections in its neural network. However, this may not be as easy as it sounds, since just like with humans, there may be a risk of interfering with other connections so that the AI may forget useful information. Also, it may not be clear which connection should be changed in the first place. So, even in the case of an AI, it may be better that all the training happens by simply inputting data—which may be carefully selected—into the system and patiently waiting until the learning happens.

### *Training can become automated*

Another major difficulty in meditation training is sustaining attention in the way typical meditation techniques require. I need to emphasize that we actually have two different attentional mechanisms at play here. First, there is *sustained* attention on the task at hand, meaning that you concentrate on meditation and don't think about anything else, as explained in Chapter 9. Second, there is *sensory, selective* attention, which means you

<sup>10</sup> *Majjhima Nikaya* 36; *Digha Nikaya* 2

<sup>11</sup> Increase in plasticity due to food restriction: (Spolidoro et al., 2011); environmental enrichment: (Sale et al., 2007); exercise and games: (Bavelier et al., 2010; Nokia et al., 2016).

<sup>12</sup> (Vetencourt et al., 2008; Castrén and Antila, 2017; Ly et al., 2018)

<sup>13</sup> (McCloskey and Cohen, 1989; Bavelier et al., 2010; Pascual-Leone et al., 2005; Kirkpatrick et al., 2017). Furthermore, too much plasticity might destroy the stability of the brain as a dynamical system, even leading to such phenomena as epileptic seizures (Kozachkov et al., 2020).

select certain data as input to sensory processing, as originally explained in Chapter 10 and extensively used earlier in this chapter. Both are necessary for successful meditation. However, sustained attention tends to be the major bottleneck because it is notoriously difficult to maintain.

In previous chapters, we have actually seen several reasons why sustained attention is difficult. First, wandering thoughts assail the mind, for example due to experience replay. But we also saw that emotions are essentially interrupts; what they are interrupting is current activity, and to do that, they have to be able to grab attention away from wherever it may be. The general concept of the brain as parallel distributed processing emphasizes the idea that there are different networks or modules which are often competing, for example, for attention and the control of attention.

Fortunately, it is possible to *learn* to use your attentional capacities better.<sup>14</sup> This is yet another form of learning, but a bit different from the typical learning we have considered: here we are talking about learning a new *skill*, as briefly described in Chapter 7. A skill means that you know how to ride a bicycle, to speak a foreign language, or to use your new smartphone; it is opposed to learning facts and increasing your knowledge about what the world is like. Skill learning follows some general laws and these apply to meditation as well. In the beginning, you need to concentrate, and spend a lot of effort, which means a lot of sustained attention. However, in time, meditation becomes more and more automated, which means that less and less conscious effort is needed. Some meditation traditions talk about meditation as “just sitting”, which is in a sense enough if the meditation is sufficiently automated. Importantly, the regulation of attention will in fact become a habit, so will be easily conducted during ordinary life, as if by itself, even outside of formal mindfulness meditation sessions.

So, there are actually two different learning processes at play: Learning that the world has certain characteristics (such as uncontrollability), and on a higher level, “learning to learn” that the world has such characteristics. The latter learning process means learning to meditate in an automated, habit-like manner, with minimum conscious effort. Thus, with practice, the meditator will be able to perform the former learning process with increasing efficiency, and this process is the one that reduces suffering according to the theory of the preceding chapter.

### *But who is actually meditating?*

The fact that meditation can become automated and habit-like means that, in a sense, it is no longer my “self” who is meditating. We find echoes of the no-self philosophy treated in Chapter 11. Some neural networks can observe the breathing without any conscious effort, or even without a conscious decision to start meditating. There is no need for any central executive to make any decision, and no need to *want* to observe the breath; it just happens. It is like when walking, you make no conscious decision to move your feet; you feel no burning desire to put one foot in front of the other.

But if the neural networks are retrained by the explicit system as I argue in this chapter, does that not mean that it is the explicit system, perhaps even a conscious self, which is in control? That might be a hasty conclusion, since there are many ways in which the control is circular. In fact, earlier (page 135) I argued that it is meaningful to say that ultimately, it is the input data that controls us. I gave the example of a meditation

<sup>14</sup>(Frieese et al., 2012; MacKenzie and Baumeister, 2015). Such learning has earlier been well-documented on a more general level in the work on self-control (Rueda et al., 2004; Baumeister et al., 2007). It can be seen as another interaction between the two systems in dual-process theory (explicit and implicit), see also (Doyon et al., 2003; Peters et al., 2011; Sun et al., 2005).

master who says that it is actually his master who is meditating, because he still hears his master's voice in his head. This shows that in order to find the “ultimate” source of control, we have to consider where the data to the explicit system comes from. Part of it clearly comes from the human society and the cultural context: there are other people that input data into us, for example in the form of meditation instructions. How that happens, and who is controlling whom, is a vast topic that I have to leave for future research.<sup>15</sup>

## Reducing interrupting desires

In addition to speeding up learning in neural networks, mindfulness meditation has further benefits. Next, we consider how it reduces suffering from the viewpoint of cognitive dynamics, which complements the frustration equation. As we saw in Chapter 13, one traditional Buddhist account on a mechanism to reduce suffering is based on the moment-to-moment cognitive chain shown in the flowchart in Fig. 13.2. The idea here is to stop the dynamic process in the flowchart in the middle so that it does not lead to its end product, which is suffering. The point where the process can best be stopped is assumed to be (in the terminology of our flowchart) the three links of desire, intention, and planning.<sup>16</sup> It is in fact assumed in early Buddhist philosophy that until the valence computation, the process is too automated, and desire provides the first link that can be stopped.<sup>17</sup>

This method is distinct from reducing desires by adopting the attitudes of the preceding chapter. Here, I am talking about sudden, “hot”, interrupting desires triggered by the valence computations, and their prevention in real-time when they are about to arise. The preceding chapter focused on reducing long-term desires from the “colder” perspective of reward calculations; this will reduce the underlying tendency for hot desires to arise, but it works only passively in the background.

One problem here is that the hot desires have the properties of interrupts, as explained in Chapter 8, which means they can be quite difficult to prevent. Therefore, it might be better to try to stop the dynamics a bit later, at the links right after desire. In Buddhism, those following links are called “attachment”, which is in our schema divided into forming an intention (i.e., committing to a goal) and planning for that goal.

Whether desire or attachment is chosen as the target, the trick here is to weaken the cognitive dynamics so that this largely automated chain leading to suffering fails to operate. If the desire or attachment is prevented from taking place, no goal is committed to or planned for, and no goal-oriented action is conducted. Thus, the whole frustration equation above is not operating, and frustration is avoided by that route.<sup>18</sup>

### *Perceptual learning*

Such stopping of the dynamics before attachment is enabled by well-known mindfulness meditation techniques. The point is to observe the cognitive dynamics repeatedly, so that one learns to introspectively detect

<sup>15</sup>My arguments here are not very rigorous since the very definition of “control” is not made explicit in this book; I simply follow typical common-sense usage of the word. A more detailed analysis would point out that control is a matter of degree: In the context of this chapter, the explicit system has only partial control of the implicit system anyway because it merely directs its attention, so the explicit system is certainly not in *total* control in any meaningful definition of the word.

<sup>16</sup>In the traditional account, these correspond to the two links of desire and attachment/clinging. See e.g. (Anālayo, 2003).

<sup>17</sup>(Mahasi, 1999, p. 89). Valence is closely related to what is called “feeling tone” or *vedanā* in Buddhist literature. I think *vedanā* can best be described as the perception of valence. (Such perception requires of course some kind of computation of valence as well.)

<sup>18</sup>Some frustration will still be felt because of the habit-based system, but as argued in Chapter 3, such frustration is much weaker than that coming from planning and execution of plans.



the different parts of the process and discriminate between the different links, in real-time. Mindfulness meditation has here the effect of training a new perceptual capacity that allows for observation of the internal mechanisms of the mind.

This is a special case of the phenomenon of “perceptual learning”.<sup>19</sup> Research on perceptual learning started in vision science by the discovery that it is possible to greatly enhance the performance in almost any visual perception task; all that is needed is sufficient training. Improvement is possible even in tasks where the limits of perception were previously thought to be set by the optics of the eye, such as the task of telling whether two lines have the same orientation (angle) or not.

In the context of meditation, such perceptual learning allows one to observe the individual elements of mental processes more accurately. An important case of such learning is that it becomes possible to observe the associations between phenomena. If B is associated with A, then, under ordinary circumstances, it may be that the thought of A immediately and necessarily brings B into mind, and it seems that A and B are two aspects of the same thing. But with mindfulness training, it is possible to see how this process breaks into pieces: First there is A, then the association is activated, and then B comes to the consciousness because of the association. This allows one to see the existence and arbitrariness of that association. In particular, one is able to dissociate desires from the stimuli that caused them, as if by creating a “space” between a stimulus (say, chocolate) and the desire, as well as the desire and the attachment that ensues.

### *Breaking the causal chain*

This opens up the possibility of breaking the long chain leading from stimulus to suffering depicted in Fig. 13.2 by learning to perceive all links in the chain more accurately, and in real-time as they are happening. Introspectively, the meditators often report that it feels as if the whole process were slowed down. By such perceptual learning, the process is also to some limited extent brought under conscious control. Even if a stimulus leads to a strong valence, the ensuing desire and the following steps will not happen completely automatically, but there is some space for deliberation. Perhaps such breaking of the causal chain is most understandable in the case of planning, which is often a rather conscious process, and as such, it should be possible to decide not to initiate it at all. Obviously, there is a strong unconscious tendency to start planning when desire arises; it is comparable to the unconscious reaction to start scratching a body part that is itching. However, with practice, such an unconscious tendency can be weakened, inhibited, and perhaps even completely removed. That would mean not letting “attachment” arise in Buddhist terminology. The key is to be able to consciously recognize when the planning is being triggered, instead of letting it happen automatically.<sup>20</sup>

It is important to achieve automatization of such mindfulness by long-term meditation practice, as described above. The learned and automated tendencies of observation can then create the possibility for inhibiting the more innate automatic tendencies of desire and attachment. In fact, if such observation is followed by conscious, deliberate inhibition of desire or attachment often enough, that very action of inhibition will become automated as well. Conscious control processes are often too slow and weak to prevent the processes underlying hot desire or other interrupts, so it is really important to train the neural networks to initiate the

---

<sup>19</sup>(Sagi, 2011)

<sup>20</sup>As already mentioned, Libet et al. (1983) proposed, rather controversially, that while the consciousness does not decide actions, it has a “veto” over actions: It can cancel an action sequence that the unconscious neural networks are trying to perform. This might provide an interesting explanation of how consciousness, in an advanced state of mindfulness and metacognition, seems to be able to prevent habitual actions (Baer, 2003; Garland et al., 2014), such as stopping the twelve-fold chain.

action of inhibition as well. Once the neural networks have been trained to perform both the detailed observation and the inhibition during formal meditation sessions, they may be able to transfer that skill to everyday life with its infinite temptations.<sup>21</sup>

## Emptying the mind and reducing simulation

Another additional benefit of meditation is that many people report feeling great pleasure when meditating. This is often attributed to the fact that the mind is strongly focused on a single object, such as breathing, and thus emptied of any thinking. Several traditional meditation schools actually maintain that an “empty” mind is happy, that is, a mind where there are no thoughts, whether wandering or intentional. (Emptiness of the mind does not here refer to the Mahayana Buddhist concept of emptiness we saw earlier.) A similar pleasurable state is sometimes achieved in the state of “flow” where wandering thoughts are equally absent.<sup>22</sup>

Understanding why an empty mind tends to be happy is one of the deepest problems for a scientific understanding of the mechanisms behind meditation, and not quite resolved at the moment. A number of viewpoints can be taken here. In a traditional Buddhist account, where desire is considered the basis for suffering, a simple explanation would be that an empty mind is happy because it does not have any desires (including aversions).<sup>23</sup> On the other hand, Chapter 9 reviewed research showing that wandering thoughts are typically related to a negative mood; however it was not clear if those results apply to all thinking and not just wandering thoughts, and what is the cause and what is the effect. Yet another viewpoint is to recall once more Cassell’s statement that “to suffer, there must be a source of thoughts about possible futures”, which cannot exist in an empty mind.

In the framework of our frustration equation, we can formulate a more computational viewpoint. Reward loss is computed every time a simulation, whether in terms of replay or planning, is conducted in the brain. A reduction of thinking should reduce mental pain since such simulation of frustration or reward loss is reduced. In fact, in our frustration equation on page 158 we have the term “how many times it is perceived or simulated” which gives the number of times the reward loss is computed. Reducing mental simulation will reduce this term, and thus suffering. Reducing mental simulation will, for example, reduce rumination over past errors, simulation of future threats to the person, as well as judgements related to self-esteem, which are some of the most important sources of suffering.<sup>24</sup>

The logic just given may explain why many meditation methods have the explicit goal of emptying the mind of thinking, or at least reducing thinking. Typically, one concentrates on a single object, such as the breath. This immediately reduces thinking, including wandering thoughts—but does not eliminate them completely,

<sup>21</sup>Related models consider how mindfulness meditation helps in addiction (Brewer et al., 2014; Garland et al., 2014). In particular, Brewer et al. (2014) proposes several mechanisms describing how mindfulness meditation can “de-automate” the dynamics, including learning to simply observe aversive states without reacting to them and taking them less “personally”, while becoming “more aware of habit-linked, minimally conscious affective states and bodily sensations”.

<sup>22</sup>(Csikszentmihalyi, 1997)

<sup>23</sup>See page 86 for a proposal on how desire and aversion in themselves produce suffering.

<sup>24</sup>A problem with this logic, which we already partly saw in Chapter 9, is that it is not clear why simulation of positive experiences would not cancel the effect of simulating negative experiences. Somehow, it seems that negative experiences are stronger in this case. It is possible that this only holds for some people whose thinking just happens to be more often negative than positive, and it is those people whose mood is most improved by meditation. Or, it could be that due to some evolutionary reasons, that is the case for the vast majority of humans: Baumeister et al. (2001) reviews a great number of results leading to the conclusion that “bad is stronger than good” as far as the emotional effects of life events are concerned. The case of rumination was treated in Chapter 9.

as the meditator soon notices. An important part of meditation is how to react to the occurrence of wandering thoughts. Some meditation techniques directly aim at suppressing them by refocusing on the original object of meditation. Suppose you have any unpleasant, possibly scary wandering thoughts about the future or the past during meditation. If you refocus on the meditation object, thus clearing the mind of such scary wandering thoughts, it is rather obvious that suffering will be reduced.<sup>25</sup> Being able to thus prevent negative wandering thoughts from occurring should have a strong positive effect on mood, in line with our logic above based on frustration equation. In fact, it has been shown that the default-mode network, largely responsible for wandering thoughts, is less activated in experienced meditators.<sup>26</sup> (Below, we will see an alternative approach to wandering thoughts based on meta-awareness.)

In Buddhist training, there is also a strong emphasis on focusing on what happens “here and now”. In other words, you learn to change your cognitive style to a more “experiential” one, which means you replace most thinking, whether future- or past-oriented, by the simple sensory experience of the present moment.<sup>27</sup> This is essentially another shift of attention away from thinking, but this time the shift is to any immediately present perceptual input, instead of any pre-selected object like the breath. It can be practised in everyday life, outside of any meditation sessions. Such an experiential cognitive style is a way of conceptualizing a long-term change in neural networks that leads to a mind which is more and more empty. It can be further motivated from the viewpoint that attentional resources are limited, and one cannot pay attention to many things at the same time. Thus, such a cognitive style can reduce attention to reward losses even in real life, and not only in simulation, since it directs attention elsewhere. In particular, any reward loss is only briefly observed without paying too much attention to it, before attention is directed to something else in the here and now. According to the frustration equation, such reduction of attention reduces suffering, this time based on the term “amount of attention paid” since reducing attention reduces the impact of any perception.<sup>28</sup>

Buddhist philosophy, as well as the theory in this book, further suggest another very different way for achieving a reduction in replay and planning, which is nothing else than adopting various philosophical attitudes described in the preceding chapter. Planning how to obtain future rewards is likely to be reduced if future rewards are considered lesser; there is simply not so much incentive anymore in planning for them. Likewise, planning to avoid threats will be reduced if those threats are seen as relatively uncontrollable. Furthermore, when the uncertainty of our thoughts and perceptions is realized, spontaneous thinking is often reduced, since there seems to be much less point in simulating something which is uncertain anyway. This is how adopting the philosophical attitudes discussed in this and the preceding chapter will also lead to a reduction in simulation, and towards an empty mind.<sup>29</sup>

---

<sup>25</sup> (Kuyken et al., 2010)

<sup>26</sup> (Brewer et al., 2011) Recapitulating some of the logic above, we arrive at a speculative computational explanation of why almost any wandering thought leads to suffering, and why the elimination of almost any wandering thoughts reduces suffering. Namely, most wandering thoughts are related to some kind of desire or aversion, which either underlies planning of future action or motivates replay of a rewarding or punishing past episode. If we combine this with the idea that aversions and desires are suffering in themselves (page 86), we see why wandering thoughts almost necessarily lead to suffering.

<sup>27</sup> (Watkins and Teasdale, 2004)

<sup>28</sup> Such an experiential style might also work by increasing (the feeling of) controllability. The past cannot be changed anymore and any control of the future is weak, while the present moment may afford much more control.

<sup>29</sup> This idea, reduction of desires leading to an empty mind, shows how the question of causality regarding emptiness of mind and happiness/suffering is complex. (See also footnote 33 in Chapter 9.) We started this section by pointing out that emptying the mind by meditation often has the effect of making people feel more joyful (as reported by many meditators), and this is consistent with the idea that the ensuing reduction of simulation should logically reduce suffering. Thus, emptying the mind was seen as an intervention

## Attitude of acceptance

There is an important caveat in any attempt to reduce mental phenomena, be it desires or wandering thoughts. It is important that this training does not lead to the idea or evaluation that the mental phenomena are bad. Such an attitude would, in itself, easily lead to aversion, and thus to suffering. In the extreme case, if there is aversion towards the mental phenomenon of aversion, that may lead to a vicious circle, which constantly increases aversion. To counter this tendency, it may be necessary to actually bring in new mental phenomena so as to neutralize the existing ones.

It may sound paradoxical to say that one should not think of the mental phenomena as bad, or at least undesirable. How could one not think that, say, desires are bad if one believes they lead to suffering? And how is one supposed to let go of them if one does not regard them as something negative, something to be avoided?

The solution to this paradox is that while the actions of the meditator should be chosen so as to reduce desires (or other mental phenomena), it is still possible to avoid creating any new aversion in the sense of a new mental process. Thus, on an abstract level, it is useful to consider the desires “bad”, or perhaps rather as something that it would be better not to have, but such thoughts should just work in the background as weakly as possible, instead of being strong and actively cultivated. In particular, they should not lead to any interrupt-like aversive emotions. Such processing is possible since the neural networks can implement automated habit-like action tendencies that try to avoid certain phenomena, and that can happen without any need to activate the desire/aversion system. As an extreme example, when you are walking, you know you that losing your balance is “bad”, but you probably don’t feel a constant aversion or fear towards stumbling; your neural networks have simply been trained to avoid that happening; they “reduce stumbling” so to say but without any aversion.

In practice, it has been found that with meditation, the tendency to develop aversion is so strong that specific techniques are necessary to reduce it. The key principle is to cultivate the attitude of acceptance. This means a general attitude of accepting all thoughts and sensations that come to the mind, instead of resisting or judging them. More precisely, acceptance here means simply not activating processes of aversion, i.e. not activating a desire to get rid of something. So, acceptance here is taken in a very limited sense; this is not about a moral acceptance, or about not thinking that some things could be bad for you. Such acceptance could also be described as removing resistance; *nonreactivity* is a related term used in current research.<sup>30</sup> For example, a depressed person may be annoyed by the very occurrence of rumination. In such a case, accepting that ru-

---

that causally reduces suffering. In contrast, the idea that reducing desires reduces (especially wandering) thoughts is in line with some classical Buddhist authors who seem to claim that the emptiness of mind is mainly an *effect* of mental development, not a cause of happiness (Williams, 2008b, p. 55). In such thinking, reducing desires reduces frustration as discussed in Chapter 14, and an empty mind is just a side-effect. Meanwhile, the discussion on the experiential cognitive style just given could probably be interpreted based on either causal direction; either an experiential style makes the mind empty, or emptying the mind leads to a more experiential style; or perhaps both are effects of the reduction of desires or some similar cause.

<sup>30</sup>Lindsay and Creswell (2017), while emphasizing the importance of acceptance in mindfulness training, use the term almost synonymously with “nonreactivity”. Hayes and Pierson (2005) define acceptance as “an open and noncontrolling stance toward all experiences”, which shows explicitly the connection to control. Meanwhile, Peacock (2018) offers an interesting interpretation of the goal of early Buddhist philosophy in terms of “freedom from enthrallment to reactive patterns”, which places nonreactivity at the very center of Buddhist training. It should be noted, however, that in actual meditation training, it is often recommended that an active, positive feeling (possibly what is called loving-kindness) is developed towards mental phenomena (Grabovac et al., 2011; Hofmann et al., 2011). It may be necessary to actively develop such positive feelings to counteract the inherent tendency to aversion and judgement; simply trying to refrain from negative judgements and practising meditation based on observation may not remove them efficiently (*Samyutta Nikaya* 10.4).

mination occurs may actually be beneficial, since it removes the suffering due to the aversion to rumination.<sup>31</sup> Again, the acceptance we are talking about has a limited meaning; one can still use various techniques to reduce the rumination.

Such an accepting attitude can actually be adopted towards all mental phenomena. In fact, many mental training systems include some kind of active acceptance practice of all mental phenomena as an integral part. Such an acceptance practice can be seen as a specific method for reducing aversions of all kinds. It complements the methods described so far, which were more oriented towards reducing desires in the restricted sense of the word (i.e., excluding aversion). It is closely related to the practice of letting go that will be considered below.<sup>32</sup>

Theories such as those explained in this book may help in the acceptance because simply understanding the mechanisms behind, say, wandering thoughts or emotions, may enable you to accept them. If you are convinced that they are natural processes which actually have some computational benefits, and that they are largely outside of conscious control, it may be easy to just let them happen, and naturally go away, without fighting against them. This is related to seeing “causality” in Buddhist terminology (considered in Chapter 14, page 162), but it goes further, since the phenomena are seen as not only natural and uncontrollable but even useful—at least from an evolutionary viewpoint. Based on this viewpoint, even frustration could be accepted as an unavoidable part of a learning process.

Ultimately, even pain and suffering need to be accepted on some level. Any aversion towards them will create a lot more suffering. As an extreme example, people suffering from chronic pain will suffer even more if they “catastrophize” the pain, resist it, and develop a particularly negative attitude towards it; accepting the pain will help.<sup>33</sup> The Buddha gave a famous simile of a man who is struck by an arrow, and thus suffers from physical pain. If he “sorrows, grieves, and laments”, feeling aversion towards pain, he makes the suffering even worse, as if he were struck by a “second arrow”.<sup>34</sup>

## Metacognition and observing the nature of mind

There is one more very particular form of attentional control operating in mindfulness meditation, especially in more advanced stages of the practice: direction of attention and awareness to a metacognitive level. Metacognition means here cognition about cognition, as in, for example, thinking about one’s own thoughts, or observing one’s own perceptual processes. In such a case, the “higher”, metacognitive part of your mind is observing the “normal” thinking or perceiving part of your mind.<sup>35</sup> Such metacognition is presumably possible because of the parallel and distributed nature of brain function, which means one part of the brain can observe what is happening in another part.

An obvious utility of such metacognition is enabling introspection that allows you to understand the processes underlying your thinking, emotions, and desires. This is of course the goal of a multitude of psychotherapeutic systems. However, the practice of meditation, especially in a Buddhist context, can go much deeper in

<sup>31</sup> (Feldman et al., 2010)

<sup>32</sup> A practical introduction to such meditation methods is provided by Brach (2004).

<sup>33</sup> (Veehof et al., 2016)

<sup>34</sup> *Samyutta Nikaya* 36.6

<sup>35</sup> (Beran et al., 2012; Proust, 2010; Fleming et al., 2012). An increase in metacognition is one of the more robust findings in studies of mindfulness training (Lao et al., 2016).

this respect, and a well developed metacognitive attitude is seen as interesting in its own right. Development of metacognition could be seen as another example of perceptual learning discussed above.

Buddhist meditative practice eventually leads to a class of advanced techniques based on *meta-awareness*, that is, the quality of the consciousness or awareness present in such metacognition. It is awareness of awareness; in other words, there is conscious recognition or perception of the fact that there is awareness. This may seem very complicated or even paradoxical, but in fact it is something that we are regularly engaged in, if only fleetingly. A typical example used in neuroscience is when you realize your mind is wandering and regain focus; that realization was on the level of meta-awareness. But there are many more interesting cases.<sup>36</sup>

Consider the following case: sensory meta-awareness. If I ask you whether you see this book, you would reply in the affirmative. I can formulate the question in a more explicit way: Are you aware of the fact that you are consciously perceiving this book? You would probably still reply in the affirmative. It is almost the same question really, since in colloquial language if you “see” something it means that you see on a conscious level. If you can consciously recognize that you are consciously perceiving this book, you must be aware of such conscious perception happening, and thus there is meta-awareness. So, you were fleetingly aware of the sensory awareness of the book; you moved to the metaconscious or meta-aware level for a few seconds. That shows that almost any kind of sensory awareness can be accompanied by meta-awareness, and it can be deliberately initiated. While this meta-awareness didn't last long, it is possible, as a meditative exercise, to stay on the level of meta-awareness for a longer period of time.<sup>37</sup>

The same kind of meta-aware observation can even be extended to thinking. Some meditation techniques emphasize observing the wandering thoughts while they are taking place, instead of suppressing them. The possibility of actually observing the wandering thoughts and their contents in real-time—instead of merely noticing that you have had some wandering thoughts a while ago—may seem quite paradoxical. However, it is possible to learn such sustained meta-awareness of one's thoughts with enough meditation practice.<sup>38</sup> From this viewpoint, at least on advanced levels of practice, it may not be necessary to reduce wandering thoughts; after all, any attempt to empty the mind may create new suffering because the mind is uncontrollable. Instead, one may change the quality of the awareness in the sense that the attention is mainly operating on the metacognitive level.<sup>39</sup> Such meta-awareness often feels like perceiving one's thoughts as if from the outside, instead of being inside or involved in them. In fact, if your mind engages in a scary simulation of something that might happen to you in the future, you can now just watch the simulation while reminding yourself that it is not actually happening, it is just a simulation where your mind plans possible courses of action. With such a quality of consciousness, there is actually little need to *stop* the simulation to reduce suffering. Moving to such a metacognitive level is actually often an automatic consequence of long practice in mindfulness training, and may easily happen during an intensive meditation retreat.<sup>40</sup>

<sup>36</sup>(Chin and Schooler, 2010; Schooler et al., 2011). As in earlier chapters, the words consciousness and awareness are here used synonymously. The distinction between metacognition and meta-awareness is not always clear and there is overlap on how these terms are used. The key difference for me is, however, that meta-cognition can happen even in an unconscious agent, while meta-awareness, by definition, requires consciousness.

<sup>37</sup>(Tejaniya, 2008, e.g. p. 77-79,121-126)

<sup>38</sup>(Tejaniya, 2008, e.g. p.126-133);(Pramote, 2013, Ch. 4);(Kyabgon, 2015, p. 177-179). The difference between intermittent (fleeting) and sustained meta-awareness in a meditation context is considered by Dunne et al. (2019). Smallwood et al. (2007) emphasizes that the absence of any metacognition, or meta-awareness, is typical of wandering thoughts.

<sup>39</sup>(Teasdale, 1999). This technique is different from but related to directing the attention elsewhere on the non-meta level, such as to the here and now that was described earlier.

<sup>40</sup>Dahl et al. (2015) discusses different meditation techniques and the role of meta-awareness in them. This is related to what is called

Now, what if you could spend a considerable proportion of your daily life on the meta-aware level? Such long-term sustained meta-awareness seems to be possible after extensive meditation practice. Importantly, such meta-awareness may lead to insights that convince the meditator about several philosophical points we have seen in this book. You may see all conscious mental phenomena, that is, all the contents of your consciousness, such as perceptions and thoughts, as results of impersonal computational processes. In other words, they are simply mental constructions, or results of a simulation performed by your brain. This logic may lead to the conclusion that even what you see in front of you at this very moment is a perception constructed by your brain, based on various unconscious inferences, sometimes hardly better than guesses: you really have no other source of information about the world but perceptions and thoughts playing in the virtual reality of consciousness. Perceptions and the ensuing thoughts are thus necessarily subjective, contextual, fuzzy, and uncertain constructs—in Buddhist philosophy, they are called empty. They do not represent any absolute truth about how things are. This is in stark contrast to our inherent tendency to think that our perceptions are somehow identical to reality.<sup>41</sup>

### *Meta-awareness and suffering*

In addition to the deep insights just described, there is another immediate utility in keeping a meta-aware attitude towards all mental phenomena: People who practice such meta-awareness often report great calm and even “bliss”. The reason is not well understood from a neuroscience viewpoint, but I would assume that less attention is paid to error signals, because attention and awareness have largely moved to the meta-level. Perhaps error signalling is somehow generally dampened, due to some mechanism to be understood. Introspectively, the effect can be described as the meditator keeping some distance to the thoughts and perceptions, and taking them less personally as well as less seriously. Going back to the frustration equation (page 158), we can assume that any reward loss will be paid less attention, meaning that meta-awareness is reducing the term “amount of attention paid to reward loss” in the frustration equation. Furthermore, the insights into uncertainty described above will further reduce reward loss by reducing the corresponding term in the frustration equation, similar to mechanisms already explained in Chapter 14.

The insights on emptiness and meta-awareness may also reduce suffering by a very different mechanism. In this book, the main approach to suffering has been to see it as frustration, and most interventions are based on that model. However, already in Chapter 2 we saw that the approach based on threats *à la* Cassell may provide an alternative in its own right. We all tend to have thoughts about bad things that might happen to us. That cloud might start pouring rain any minute; that car might run over me, and so on. Such simulations about what might go wrong in the future are essential for the particular kind of suffering based on threats to the intactness of the person, as emphasized in Cassell’s definition (see page 17). Now, using the concepts of emptiness and meta-awareness, we can develop further interventions against suffering based on threats.

---

*decentering* by Fresco et al. (2007); Safran and Segal (1990).

<sup>41</sup> This eventually leads to what is the deepest form of meta-awareness: being simply aware of the existence of consciousness itself, or of the very capacity to be conscious of mental phenomena, as opposed to being aware that you are aware of some specific phenomenon (such as the perception of this book in the example above). Such awareness is related to what is called seeing the (true) nature of mind (or consciousness) in Buddhist and related literature (Brahm, 2006; Dalai Lama et al., 2011; Kyabgon, 2015; Spira, 2017); it is also related to the attitude briefly described at the very end of Chapter 12. However, it is outside of the scope of this book, and goes well beyond our frustration equation. (Nevertheless, similar to what is argued next in the main text, it could lead to minimal attention to reward loss and therefore a great reduction of frustration even according to that equation.)

Earlier we already saw that meditation can enable learning that an association triggering fear is no longer valid, thus leading to fear extinction. But suppose you were able to see all threats as empty: uncertain, subjective, open to interpretation, nothing but mental constructs. This is particularly feasible in case of threats to your self-esteem or social status instead of your survival; or if the threat only happens in a simulation. Then, any threats would not be taken that seriously, and their ability to trigger fear would be weakened. From this viewpoint, it is not necessary to look at the desires, such as self-needs, that underlie the threats as we did in Chapters 6 and 14. It is now possible to directly intervene on the threats themselves by seeing them as empty. From a historical viewpoint, this suggests an interpretation of the Mahayana school complementing the Buddha's original frustration-based methods by offering interventions that more directly apply to threats.<sup>42</sup>

## Letting go and relaxation as unifying principles

Buddhist philosophers often use the concept of “letting go” to recapitulate the general attitude that has been described in this chapter and the preceding one. At the most concrete level, the idea is that we let go of things and objects in the sense that we don't strive to possess or control them anymore. On a more computational level it means we let go of desire, i.e. we don't even want those things in the first place—nor do we want to avoid them. The same approach can further be applied to thoughts and perceptions, which are understood to be subjective and unreliable, so they can be let go of. Feelings and emotions are likewise just observed and then let go of. This whole simulation is not taken that seriously anymore. Combined with the no-self philosophy, the attitude can be recapitulated as letting go of everything which is not part of me, and since nothing really is part of me, or my “self”, everything is let go of.<sup>43</sup>

Letting go is an expression that obviously has a clear connection to the term *reduction* we have used very often, especially in the preceding chapter, even in many section titles. It is not so much a question of programming new routines or new functionalities. The idea is to reduce activity, letting go of existing mental associations and routines. The key is *less* desire and aversion, less replay and planning, fewer interrupts, and so on.<sup>44</sup>

An important point about letting go is that it circumvents the paradox of wanting not to want anything. If meditators want to reduce of desires, they can be seen as wanting not to want, which may sound impossible. This apparent paradox in Buddhist philosophy has been pointed out by a number of authors: since wanting

<sup>42</sup>The Buddha may not have talked very much about fear in his discourses, with the exception of the fear of death (e.g. *Anguttara Nikaya* 4.184); but see footnote 31 in Chapter 12 for a Theravadan quote focused on fear. I would further think that something similar to the cognitive dynamics in the flowchart in Fig. 13.2 may apply to threats directly, when desire is replaced by threat and the ensuing fear. Furthermore, the experiential style of focusing on the “here and now” is has an effect similar to meta-awareness, since with that cognitive style, future consequences of threats will not be simulated at all. The connection between desires and threats is highly complex, and an interesting topic for future research.

<sup>43</sup>To quote *Samyutta Nikaya* 35.101: “Whatever is not yours: let go of it. Your letting go of it will be for your long-term happiness and benefit. And what is not yours? The eye is not yours: let go of it. (...) [Visual] forms are not yours: Let go of them. (...) Eye-consciousness [i.e. visual awareness] is not yours: Let go of it. [The text goes through all the sensory organs, the objects of sensation, and the accompanying sensory awarenesses.] The intellect is not yours: let go of it. (...) Ideas are not yours: let go of them. (...) Whatever arises (...), experienced either as pleasure, as pain, or as neither-pleasure-nor-pain, that too is not yours: let go of it. Your letting go of it will be for your long-term happiness and benefit.” (Translated by Thanissaro Bhikkhu)

<sup>44</sup>Alternatively, letting go could be seen as the opposite of attachment, especially if the corresponding term (*upādāna*) is translated as “grasping” or “clinging”. However that would require an interpretation of attachment which is quite different from what I have done in this book.



not to want is a form of wanting, how could one possibly get rid of wanting by such wanting? The paradox is actually so obvious that even the Buddha himself, as well as his immediate disciples, were confronted with claims that his system is inherently paradoxical. Thinking of the mental process of Buddhist training as letting go, and as reduction, should largely resolve this paradox of seemingly wanting not to want. Letting go conveys a reprogramming that reduces mental activity instead of introducing a new desire.<sup>45</sup>

One way of interpreting letting go is that it is mental *relaxation*, in the intuitive sense of absence of activity and tension. Desire, and the subsequent goal-setting, are actively engaging in a mental activity, and thus they are a kind of opposite to relaxation. Figuratively speaking, just as muscular activity prevents relaxation, wanting is opposite to relaxation in that it relies on specifically activating certain computational processes. If you set the goal that you don't want anything, you would actually be just setting one more goal, and increasing mental activity—this is another viewpoint to the paradox we just saw. But if instead, you learn to relax the planning and goal-setting system so that it simply rests, and does not set any goals and does not plan, then you resolve the paradox of wanting to not want. Learning such relaxation is not easy, but the training methods discussed in this book were basically all designed to lead towards such a mental relaxation.<sup>46</sup>

The ultimate goal of Buddhist training is called *nibbāna* or *nirvāṇa*, depending on which ancient Indian language is used. It is defined as a state devoid of any suffering, the cessation of all suffering. The term literally means *extinction*, as in a fire being blown out. It is often described in negative terms such as “unconditioned”, “unconstructed”, or even “unborn”, which may sound nonsensical. I think the key to understanding this is that *nibbāna* is reached by reducing, and ultimately removing, various mental phenomena, in particular desire; it is not about constructing any new mental phenomena. This may again sound paradoxical to any beginning meditator struggling to maintain even a tiny amount of concentration, but I am of course talking about highly advanced stages of practice here.<sup>47</sup> Thus, the best description of the ultimate state may be entirely negative, in terms of what it is *not*, and what it does *not* contain. It is often described as freedom, and in particular it is freedom *from* those elements of the mind that produce suffering.<sup>48</sup>

---

<sup>45</sup>A detailed resolution of this paradox also needs to look at how the meditation practice changes in time, over a time span of many years. Initially, meditation is based on the desire to reduce suffering, and makes use of the desire to let go. But ultimately, you let go of even the desire to be happy, and, paradoxically, of the desire to let go. This is not a contradiction since you let go of letting go only after a long practice, so you have let go of other things already. The new attitudes and habits required for reducing suffering are now automated in your neural networks and need no effort or explicit desire to operate anymore. Note that this is clearly related to the problem of “aversion towards aversion” that we considered above in connection with acceptance. In a similar vein, Striker (2004) emphasizes that Pyrrhonian Sceptics did not (actively and purposefully) suspend judgement, as it is sometimes claimed, but rather were unable to arrive at any judgement and gave up any such attempt. For a traditional account of how one of the closest disciples of the Buddha replied regarding this paradox, see *Samyutta Nikaya* 51.15; for a modern philosophical approach that emphasizes letting go, see Herman (1979).

<sup>46</sup>Thus, in the end, we find another paradoxical process, “relaxing the desire to relax”. I would speculate it is a bit like going to sleep. You create the right conditions by lying down, and in some sense, emptying your mind. But you cannot simply decide to fall asleep. In fact, having a strong desire to fall asleep will make it more difficult. You will fall asleep naturally if the right environmental conditions are met, and you have the right state of mind. Relaxation of the mind as described in the main text is a slightly less natural process, it seems, and you may actually need to practise various techniques to create the right conditions. But it is possible to relax without having a strong, active desire to relax, and indeed, to obtain total relaxation, you need to even relax the desire to relax.

<sup>47</sup>Mahasi (2016) gives a traditional Theravadan commentary: “Because there is no arising in the *nibbāna* element [which is the cessation of conditioned phenomena through their non-arising], it is called not-born (*ajata*) and not-brought-to-being (*abhhuta*). Because it is not made by a cause, it is called not-made (*akata*). Because it is not made dependent on causes and conditions, it is called not-conditioned.” (see his Chapter “Attainment of Fruition”).

<sup>48</sup>While the Four Noble Truths indicate that extinguishing desire accomplishes the goal of removing suffering, aversion (or hate) and ignorance (or delusion) are usually added to the list of phenomena that have to be extinguished, see e.g. *Samyutta Nikaya* 38.1. (The

One might think such a mind-state with no contents must have neutral valence, and could even be boring. Yet, Buddhist philosophy claims it is extremely happy and pleasant, in fact pure bliss. It is claimed to be the only thing that is not unsatisfactory in any way. This may perhaps be understood if we consider the mind in such a state to be *completely* empty, and we have seen that even a relatively empty mind seems to be, for some reason, quite happy.<sup>49</sup> Nevertheless, we find yet another interesting paradox: How can having a completely empty mind possibly be pleasant, since it logically should not contain any pleasure either. I will not try to resolve this paradox which seems to reach metaphysical depths; let me just quote Sāriputta, one of the closest disciples of the Buddha, who put it very simply:<sup>50</sup>

Just that is the pleasure here, my friend: where there is nothing felt.

---

exact meaning of ignorance/delusion in this context is quite controversial.) Such lists come in various lengths, and ultimately may contain almost all mental phenomena, as when the Buddha says that he teaches “for the elimination of all standpoints, decisions, obsessions, adherences, and underlying tendencies, for the stilling of all formations, for the relinquishing of all attachments, for the destruction of craving, for dispassion, for cessation, for Nibbana.” (*Majjhima Nikaya* 22). It should also be noted that the conception of nibbāna or nirvana is quite variable among different Buddhist schools. For a detailed account of the early Buddhist view, see Harvey (1995). In later Buddhism, there is more emphasis on the extinction of conceptual thinking—as when Nagarjuna says that nirvāna is “the calming of all verbal differentiations” (Williams, 2008b, p. 75)—and realization of the “nature of mind” (Kyabgon, 2015, e.g., p. 156) which is an advanced form of meta-awareness.

<sup>49</sup>In particular, the mind might be empty of all perception in addition to thinking, even including proprioception and interoception (feeling of the body, see footnote 24 in Chapter 12), which are partly the basis of the feeling of “self”. The Japanese Zen master Dogen said that he experienced the “dropping away of body and mind”, while Brahm (2006, p. 158) emphasizes that in deep meditative absorption (jhāna), “the five senses have shut down”. Clearly, such complete emptiness can only be achieved by letting go of everything, not by making an effort to empty the mind.

<sup>50</sup>*Anguttara Nikaya* 9.34

## Chapter 16

# Epilogue

There is a wide consensus that trying to build an AI teaches us a lot about what human intelligence is about: an AI works as a model of the human mind. I think it is the same for suffering. For sure, a model is not the same as the real thing; some things are always missing. You cannot actually drive to work with a computational model of a car; mathematical equations of physical forces and chemical reactions written on a piece of paper do not actually make your car accelerate. Yet, it is such models that enable the construction of cars and even rockets that fly to the moon.

A good model can tell us a great deal about the real thing, and thus help science understand how a complex system works. A model can also enable us to predict what the system does in the future, for example, by providing a weather forecast. But from the viewpoint of this book, what really matters is if the model is predictive in the following narrow sense: Does it enable us to predict what results *interventions* have on the system? That is, does it help us in changing the system in some way we find preferable?

This book proposes that computational models of human suffering can tell us what kind of processes are *necessary* for suffering. The AI models in this book explicitly showed us some of the conditions, causes, and processes that have to be operating in order that suffering arises. That means we can develop methods that will *reduce* suffering: We simply need to remove the necessary conditions, or, at least, make them weaker. This is why I think the models in this book are useful, and the later chapters of this book were, in fact, all about methods to reduce suffering.

It is possible to argue that an AI or a robot cannot *really* suffer since it is not conscious. In other words, the computational processes considered in this book may not be *sufficient* for suffering if one insists that suffering must be conscious. However, that is beside the point if our main goal is to develop methods that reduce suffering. Actually, some even claim an AI is not really intelligent—according to some stringent conditions for intelligence—yet AI is not only capable of performing some very useful practical tasks, but it has also greatly advanced human neuroscience by giving insight to the computations performed by the brain.

The interventions I proposed were mostly identical to what existing philosophical systems propose, while I showed how to motivate them using current AI theories. The theory in this book will hopefully be complemented by further research; I think this is just the very beginning of a long-term scientific endeavour. I hope it will lead to more and more efficient interventions in the future, including completely new kinds of interventions.

I certainly do not claim that the theory in this book would be either complete or perfect. In particular, there are quite probably mechanisms of suffering which do not fit into the framework of this book. That may

be the case, for example, for suffering due to certain kinds of social emotions, or existential suffering such as lack of meaning of life. The theory in this book also attempts to explain all kinds of suffering—including self-needs, uncertainty, uncontrollability, negative emotions (such as fear and disgust), and stress—by the single mechanism of frustration. Whether such a theory based on a single mechanism is satisfactory remains to be seen in future research. For example, some interpretations of Buddhist philosophy further maintain that desire and aversion *in themselves* are suffering, and it is not quite clear how that fits the framework in this book.<sup>1</sup> As always in science, the theories could even be rejected, at least partly, as science progresses.

*Summary: Limitations of the agent lead to errors and suffering*

To recapitulate the book in a few paragraphs: we saw several ways in which the limitations of the agent and its intelligence lead to suffering. We can succinctly summarize the main problems as uncontrollability, uncertainty (or unpredictability, or impermanence), and unsatisfactoriness (including insatiability and evolutionary obsessions). The agent cannot control its environment as much as it would like; it is not able to perceive or predict the world with much certainty; it strives endlessly at questionable goals which in humans are given by evolution. Such an agent can never find satisfaction.

We saw that the inability to control the environment is partly due to the limitations of the agent's physical body and any other means it may use to manipulate the environment. Yet, to a large extent, it is also due to the agent's limited computational capacities and limited data. If there has not been enough data to learn from, the agent cannot have a good model of the world; that is, it cannot understand the regularities of the world. Even if the agent had a huge amount of data to learn from, it cannot learn well if it does not have sufficient computational capacities. Furthermore, even if the agent had a good world model, it may not have enough computational capacities to use the model properly when choosing actions, especially when planning action sequences. Likewise, perception is limited by deficiencies in the sensory input and the ensuing inverse problem. These are some of the problems that an AI, a human, an animal, and in fact any sophisticated cognitive system will encounter.

The brain, having a particularly sophisticated cognitive architecture, uses some clever tricks to try to cope with some of these problems. Wandering thoughts speed up learning by running learning algorithms in the background; however, they make us experience simulated suffering in addition to the real one. Emotional interrupts are useful when unexpected things happen and the computational resources need to be redirected, but they can be mistuned and lead to unnecessary alerts and suffering. Highly intelligent agents may have to use parallel and distributed processing where it is no longer clear if anybody is in control; such uncontrollability of the mind itself is reflected precisely in the emotional interrupts and wandering thoughts.

It is due to these limitations—physical and cognitive—that the cognitive system will make errors in its predictions, its plans, and its actions. We saw that suffering is basically a function of the constant evaluation that an intelligent system performs regarding its actions, resulting in an error signal. Without such evaluations, the performance cannot be improved. In particular, error signalling is necessary for the system to learn and update its model of the world. Yet, the constant monitoring and signalling for errors creates constant suffering. In animals and humans, we also find processes related to self-preservation and self-evaluation, which create another layer of suffering. This is what leads to the simple maxim in the title, *intelligence is painful*.

---

<sup>1</sup>See Chapter 7, page 86 for discussion on this point.

*Does intelligence necessarily lead to suffering?*

It could thus be argued that suffering is the price to pay for intelligence: without some kind of error monitoring, learning is not possible. It is common sense that errors due to past decisions have to be detected to learn to make wiser decisions in the future. Error signals might not be needed if the agent were programmed to be sufficiently intelligent to begin with, so that it would not need any kind of learning, but current AI research suggests that intelligence without learning is very difficult to achieve.

Yet, one might ask if the price is too high, whether intelligence is worth the suffering.<sup>2</sup> Would you prefer to be a bit dumber if that reduced your suffering? Suppose a drug were developed which abolishes any error-signalling in humans; perhaps that is possible by interfering with the dopamine metabolism. Suppose that as a logical side effect, it prevents you from learning new reward associations. Would it be worth taking? Actually, we don't even need to consider such an extreme case where all error signals are removed. How about just taking a small dose of that drug, so that error-signalling is reduced to some extent? You would suffer less but perhaps learn new things a bit more slowly. What would be the right balance between maximizing performance and reducing suffering?

I argued earlier that many desires are actually not good for us, and should be seen as evolutionary obsessions. Some desires are insatiable, so trying to learn to satiate them is a fool's errand. And, perhaps most importantly, a large proportion of desires are just impossible to satisfy due to the uncontrollability of the world. Clearly, frustration in those cases should be avoided altogether; they present no real trade-off between suffering and intelligence. If you really want to be frustrated, better do it in cases where the desires actually serve a useful purpose, and you learn to act more efficiently in a meaningful context.

On the other hand, even if we admit that a certain amount of suffering is necessary as a trade-off to achieve intelligence, is it really necessary that such error-signalling should be *consciously* experienced as suffering? Even the most rudimentary AI computes errors while hardly being conscious. We would not say that a thermostat, arguably the simplest possible system with some intelligence, is suffering or feels pain when it realizes the temperature of the room is not what it is supposed to be. This leads to another thought experiment: How about a drug that does not reduce error-signalling, but prevents it from reaching our conscious perceptions—would you not take it? In fact, this need not be just a thought experiment. Moving to the level of meta-awareness, as described in Chapter 15, seems to reduce the felt impact of all suffering, a bit like such a drug.

Consciousness is a great mystery. It cannot be entirely avoided in any discussion on intelligence or suffering, but unfortunately, there is very little we can say with any certainty. One thing which is clear, though, is that the way human consciousness usually operates is not very nice from the viewpoint of suffering. A large amount of suffering is even created out of nowhere by conscious simulation.

*From intelligence to wisdom*

Nevertheless, intelligence may not necessarily be only a bad thing from the viewpoint of suffering. Intelligence may lead to reduction of suffering once it reaches a certain stage, while being embedded in a culture that actively questions where suffering comes from and what can be done. Buddhist philosophy, together with the Stoics and other related systems, proposes that we should adopt certain ways of thinking which counteract, and to some extent neutralize, the causes of suffering. For example, we should give up any attempts to control

---

<sup>2</sup>This question was already considered at the end of Chapter 14 from a slightly different angle.

and accept that things are just happening; we should recognize that actually we don't know much and are always making decisions under uncertainty; we should give up the meaningless and even destructive desires programmed by evolution. Ultimately, we should recognize the true nature of our consciousness, that we are operating in a kind of virtual reality, which bears only some indirect relation to the actual reality.

Such proposals are quite radical, and have been recognized as such for centuries. This is not surprising since reducing desires and giving up control are strictly against our evolutionary programming. However, it may not be necessary to follow these ideas to any extreme extent: Buddhist philosophy in itself proposes the “middle way”, the idea being that going to any extremes is, in the end, counterproductive. Instead of giving up all control, for example, we might just give up some of the control, preferably on those things where claiming control is most clearly conducive to suffering.

Most philosophical systems that discourage acting out our desires do recognize that a human being needs to take some actions; they do not recommend complete inactivity as some might assume. Stoicism as well as Taoism emphasize acting “naturally” (or according to one's nature), which I would interpret in terms of the habit-based, automated action selection: learned associations between the current state and actions may still remain even if no reward is expected or predicted anymore.<sup>3</sup> In early Buddhist thought, motivation for mental development is often seen to be a desire of a special kind that should not be eradicated, thus providing another motivation not based on reward maximization.<sup>4</sup> Some parts of Hindu philosophy suggest performing one's duty without any concern for reward,<sup>5</sup> while later Buddhist philosophy<sup>6</sup> emphasizes altruistic action as the ultimate motivation for fully enlightened beings.

Indeed, this book has almost completely neglected the social aspects of being human—perhaps because AI's are not very social at the moment, and relevant computational theory is scarce. The theory of this book is clearly applicable to the social domain in the sense that social interaction creates its own input data that the agent can learn from. The agent might then realize that other agents are often quite unpredictable and uncontrollable, and that this leads to a lot of frustration. Yet, social interaction creates completely new phenomena which are outside of the theory of this book, but should to be considered to see the whole picture of human suffering.

---

<sup>3</sup>Epictetus recommends to “behave conformably to nature in reaction to how things appear” (*The Enchiridion*, Paragraph 6), see also the discussion on Marcus Aurelius by Hadot (2002). Laoze (Laozi) recommends “nonaction” or “effortless action”—a highly complex concept with many interpretations—which can be interpreted as acting naturally, or rather, in an automated way in our terminology. According to Chan (2018), it “seems to be used more broadly as a contrast against any form of action characterized by self-serving desire”... “nonaction would be ‘normal’ action in the pristine order of nature, in which the mind is at peace, free from the incessant stirring of desire.” Alternatively, such acting naturally could mean that any attempt to control is minimized by choosing courses of action which are in harmony with the environment; this interpretation does not, however, explain where the ultimate motivation for action comes from. In any case, it seems important that such natural action is still constrained by sound moral principles, so that it does not mean just doing whatever one feels like.

<sup>4</sup>In early Buddhist philosophy, desires for spiritual development and similar things are called *chanda*, often translated as “aspiration”. However, I'm not aware of any principled way of distinguishing between *chanda* and the “bad” desire (called *tanhā*), so this seems to be just assuming an arbitrary exception to the general theory. In fact, in later Buddhist philosophy, even getting rid of the desire for spiritual development is considered important, in line with the discussion on letting go at the end of Chapter 15.

<sup>5</sup>*Bhagavadgita* 2:47 says “You have a right to perform your prescribed duty, but you are not entitled to the fruits of action”. Duty is a socially defined concept, and as such, outside of the scope of this book. Stoic philosophy can be seen in this light as well, if the Greek *kathekon* is translated as “duty”, see p. 172 in Hadot (2002). Acting according to God's will is another formulation used by Epictetus (e.g. *The Discourses*, IV:1) and of course by various religious systems.

<sup>6</sup>For the Mahayana school, see Oldmeadow (1997), but similar ideas can certainly be found in Theravadan school as well (Brahm, 2006, p. 245); see also later in the text.

It can be argued that social interaction is essential for understanding what it is like to be human.<sup>7</sup> One aspect is that human philosophical systems considered in this book are products of a long cultural evolution. It is difficult to see how any single AI could conclude, by itself, that desire produces suffering (or errors) and should be reduced. It is probably impossible for even any single human being to discover anything like those aforementioned philosophical systems. What is necessary is a cultural learning process based on sharing information between individuals, eventually leading to accumulation of knowledge over many generations.<sup>8</sup> Such culturally produced, higher kind of intelligence, which can even consider the very concepts of intelligence and suffering as the objects of its analysis, is close to what would better be called *wisdom*. It is something much deeper than intelligence, and presumably unique to humans.

### *From individual desires to altruism*

Another essential aspect of social interaction is the human capacity for compassion, love, and similar social emotions. In classic Buddhist training, there is a group of practices based on cultivation of positive social, interpersonal emotions such as compassion and “loving-kindness”.<sup>9</sup> Interestingly, such emotions can even be directed towards oneself: As an important example, self-compassion, i.e. compassion directed towards oneself, may strongly reduce negative self-evaluations, and thus self-related suffering.<sup>10</sup> Another book could possibly be written where reduction of suffering is approached from the viewpoint of such positive social emotions. Unfortunately, any related computational theory is rather lacking at this moment.

Historically, within Buddhism, a self-centered approach to reducing suffering was increasingly criticized in the centuries after the Buddha’s death. Consequently, the later Mahayana schools adopted unselfish behaviour as the ultimate ideal, instead of your individual nirvāṇa. They proposed that it is better to sacrifice one’s own bliss and meditation time, at least to some extent, in order to help others to reduce their suffering. Slightly paradoxically, such a prosocial attitude is then seen as leading to an even higher form of happiness. I would assume that such enlightened altruistic action somehow avoids the frustration process, perhaps because there is no longer any consideration for rewards that the agent itself will get, so in a sense, the self-based desire is no longer operating. It also seems that altruistic action gives its own evolutionary rewards,<sup>11</sup> and can even provide meaning to one’s existence.<sup>12</sup> Thus, altruistic action, if performed with the proper attitude, may be the ultimate exercise to reduce suffering—even to the very person performing the action. To conclude, let me quote the Mahayana Buddhist philosopher Śāntideva who recapitulates this brilliantly:<sup>13</sup>

All those who suffer in the world do so because of their desire for their own happiness.

All those happy in the world are so because of their desire for the happiness of others.

---

<sup>7</sup> (Hari et al., 2015)

<sup>8</sup> It is crucial that the shared knowledge is cumulative, i.e., increases from one generation to another, which seems to be extremely rare with animals. A suprisingly important mechanism in such cultural learning seems to be imitation, although it might seem very primitive and unrelated to any higher form of intelligence (Iacoboni, 2005; Whiten et al., 2009).

<sup>9</sup> For current research, see Graser and Stangier (2018); Hofmann et al. (2011); Cassell (2002), and for the related emotions of forgiveness and gratitude, see McCullough and vanOyen-Witvliet (2002); Emmons and Shelton (2002). For practical meditation guidance, see e.g. Salzberg (2002).

<sup>10</sup> (Neff et al., 2007)

<sup>11</sup> For evolutionary theories of altruism, see Wright (1994); Nowak et al. (2010)

<sup>12</sup> On meaning in life and its relation to happiness, see Baumeister and Vohs (2002); Martela (2020).

<sup>13</sup> Śāntideva’s *Bodhicaryāvatāra*, written in the 8th century CE, translated by Kate Crosby and Andrew Skilton, OUP, 1995. Such an altruistic attitude is often called the bodhisattva ideal in Buddhist literature (Garfield, 2010; Williams, 2008b).

# Index

- acceptance, 188
- action-value, 50
- addiction, 94
- African savannah, 163
- agency, 124
- agent
  - definition, 22
- alarms, 97
- AlphaGo, 78, 92, 103
- altruism, 58, 198, 199
- amygdala, 73
- analog, 43
- anger, 94
- Aristotle, 18, 43, 45, 69
- associations
  - unconscious, 75, 177
- attachment (Buddhist), 85, 156
  - vs letting go, 192
- attention
  - bottom-up, 91
  - different kinds, 120
  - selective, 119
    - in meditation, 182
  - sustained, 99
    - in meditation, 182
  - to reward loss, 158, 187, 191
- automatization, 78
  - of meditation, 182, 185
- aversion, 165
  - definition, 28
  - to desires, 188
- awareness, *see* consciousness
- Bayesian, *see* prior, Bayesian
- belief-desire-intention theory, 29
- Bellman equation, 50
- binary values, 43
- bodhisattva, 199
- body, 109
- broadcasting, 20, 21, 91, 109
- Buddha, 19, 68, 125, 162, 171, 177, 189
- categories, 43, 79
- causality (Buddhist), 162
- central executive, 129
- clinging, *see* attachment
- cognitive style, 187
- compassion, 199
- conditioning
  - classical, 56
  - extinction, 181
  - instrumental, 51
- consciousness, 136
  - and pain, 15
  - and self, 144
  - brain basis, 141
  - experience, *see* subjective experience
  - hard problem, 140
  - in AI, 142
  - in animals, 141
  - meta-, 190
  - necessary for pain and suffering, *see* subjective experience
  - phenomenal, 137
  - utility, 137



- contemplation, 177
- context-dependence
  - of categories, 83
  - of learning in neural networks, 76
  - of perception, 121
- control, 124
  - as illusion, 133
  - as percept, 132
  - by input data, 135, 183
  - lack of, *see* uncontrollability
- counterfactual regret, 52
- creativity
  - positive emotions, 96
  - wandering thoughts, 107
- curiosity, 66
- data structures, 79
- death, 63, 163
  - contemplation of, 172
- decentering, 190
- Deep Blue, 30, 49
- default-mode network, 100
  - in meditation, 187
  - in rats, 106
- demon
  - Descartes's evil, 146
  - Laplace's, 166
- depression, 62, 108, 173, 188
- Descartes, 144, 146
- desire
  - and self, 65, 68, 199
  - as elaborated intrusion, 93
  - as interrupt and hot, 92, 184
  - definition, 27, 93
  - occurrent/interrupting, 93
  - reducing, 184
  - to have no desire, 192
  - wanting vs liking, 55
- digital, 43
- discounting, 48
- disgust, 92
- distributed processing, 33, 119
  - and self, 124
  - and uncontrollability, 126
  - necessity in computers, 129
- dopamine, 55, 59, 86
- dreaming, 102
- drugs
  - cocaine, 55
  - increasing plasticity, 182
  - producing misperception, 122
- dual process, 72, 118
  - communication between, 177, 180
- dual system, *see* dual process
- dukkha, 152, 164
- dynamics
  - Buddhist model, 156
  - cognitive, of desire, 154, 184
  - intrinsic, 100
- effort
  - and automatization, 79
  - and dual-process, 72
  - and perceived reward, 52
  - heuristic, 121
  - in planning, 49
  - in vision, 111
- elaborated-intrusion theory, 93
- emergence, 45
- emotions, 89
  - altruistic, *see* altruism
  - as hard-wired programs, 94
  - as interrupts, 91
  - Frijda's theory, 94
  - moral, 95
  - positive, 96
  - social, positive, 199
  - somatic markers, 98
- emptiness (Buddhist)
  - definitions, 168
  - in flowchart, 152
  - of all phenomena, 191
- empty mind, 109, 186, 194
- energy consumption, 127

Epictetus, 19, 126, 162, 164, 169, 171–173, 198

Epicurus, 171

error

classification e., 38

prediction e., 38, 55, 56

reward prediction e., 53, 57

mathematical definition, 53

error signalling, 20, 65, 96, 109, 152, 153

broadcasting, *see* broadcasting

general, 86

in control, 132

why conscious, 197

ethics, *see* moral

evolution, 10, 57

African savannah, 58, 96

and happiness, 58, 107, 109

as optimization, 40

shortcuts, 21, 143

vs learning, 57

exaptation, 80

expectation

definition, 52, 86

of reward, 122, 161

reducing, 180

experience

subjective, *see* subjective experience

experience replay, *see* replay

experience sampling, 100

expert system, 44

extinction

as in nibbāna/nirvāna, 193

of classical conditioning, 181

fear, 69

and prediction, 71

as frustration, 70

as interrupt, 92

dual process view, 73

evolution, 96

extinction, 181

phobia, *see* phobia

feedback control, 132

flow, 186

and interrupts, 96

forgiveness, 199

four noble truths (Buddhist), 19, 171, 193

free will, 133

Freud, 74

frustration

and general errors, 86

as reward loss, 17, 52

as reward prediction error, 53

as threat, 84

based on desire, 27

based on desire vs expectation, 86

based on planning, 26

based on predictions only, 54

definition of baseline, 56

definition of expectation, 52, 86

equation, 158

of internal needs, 68

of self-needs, 65, 96, 153

root causes, 87

self-needs, 85

summarizing different aspects, 85, 158

the f. equation, 158

fuzzy, 81, 169

game theory, 95

games

backgammon, 103

checkers, 103

chess, 30

Dota2, 103

Go, 78

generalization, 83

GOF AI, 42, 74, 79

Good old-fashioned AI, *see* GOF AI

gradient descent, 39

stochastic, 39, 42, 75, 77, 128, 178

gratitude, 172, 175, 199

greed, 60

habits, 51

combining with planning, 77

- slow to learn, 77
  - vs emotions, 95
- Hebb's rule, *see* learning, Hebbian
- Helmholtz, 114
- heuristics, 29
  - availability h., 164
  - dual-process view, 78
  - effort h., 121
  - somatic markers, 98
  - state-values, 49
- homeostasis, 16
- Hume, 58, 99, 123, 144
- IASP, *see* pain, IASP definition
- idealism, 146
- ignorance (Buddhist), 156, 173
- illusion, 173
- impermanence, 151, 163, 166
  - as nonstationarity, 163
  - as uncertainty, 152
- independent component analysis, 101, 116, 117
- individual differences, 97, 108, 121, 178
  - in effects of thought wandering, 186
- inductive bias, 151
- inference
  - Bayesian, 114
  - unconscious, 113
- information processing
  - analog vs digital, 43
  - distributed, *see* distributed processing
  - parallel, *see* parallel processing
- insatiability, 59, 152, 164
- intactness of the person, 17, 64, 65, 69, 84, 191
- intelligence
  - definition, 32, 76
- intention
  - and attachment, 85, 156, 184
  - definition, 28
  - reducing, 172, 184
- interpretability, 80
- interrupt theory, 91, 130, 153
- intervention, 153, 195
- invariance, 117
- inverse problem, 112
- James, William, 65
- language, 80
- Laozi, 198
- Laplace, 114, 166
- learning
  - and GOFAL, 44
  - as minimization of errors, 38
  - context-dependent, 76
  - definition with neural networks, 35
  - distributed, 131
  - federated, 131
  - Hebbian, 42, 56, 75, 115, 116, 177, 178
    - three-factor, 51
  - imitation, 78, 199
  - in the brain, 35, 116
  - iterative, 39, 75, 103
  - of Bayesian prior, 115
  - perceptual, 185, 190
  - reinforcement, *see* reinforcement learning
  - self-supervised, 117
  - skills, 79
    - in meditation, 182, 185
  - slow in neural networks, 75, 178
  - supervised, 41
  - to learn, 62, 183
  - to solve planning, 49
  - transfer, 76
  - unsupervised, *see* unsupervised learning
  - vs evolution, 57, 115
- letting go, 192
- local minima, 67
- loving-kindness, 188, 199
- Madhyamaka, 168
- meditation, 178
- mental pain, *see* suffering
- mental simulation, *see* simulation
- metacognition, 189
- middle way, 174, 198

- mind wandering, *see* wandering thoughts
- mindfulness, *see* meditation
- model, 195
- Montaigne, 71, 109
- Moore's law, 127
- moral, 63, 92, 95, 134
  - and animal consciousness, 141
  - and evolution, 58
  - expectation, 52
- neural networks
  - definition, 35
  - learning, *see* learning
- neuron, 32, 33
- nibbāna/nirvāna, 193, 199
- no-self, 151
  - as no central executive, 131
  - Buddhism, 125, 134, 162
  - deconstructing self category, 171
  - Hume, 144
  - in meditation, 183
  - letting go of self, 192
  - neuroscience, 134
  - no-doer, 134
  - realization in meditation, 180
  - vs Descartes, 144
- nonreactivity, 188
- nonstationarity, 152, 163
- objective function, 40
  - based on errors, *see* error
  - based on rewards, 48
- obsessions
  - evolutionary, 59, 152, 165, 170
- optimization, 38
- overgeneralization, 83
- pain
  - and consciousness, 15
  - as interrupt, 92
  - chronic, 58, 71, 189
  - evolutionary definition, 16
  - IASP definition, 15, 64, 65
  - uncertainty/controllability, 69
- parallel processing, 33, 113, 119, 126, 128
  - in vision, 128
  - necessity in computers, 126
- pattern recognition, *see* vision
- perception, *see also* vision
  - context-dependence, 121
  - subjectivity, 120
- phobia, 84, 97
- planning, 102
  - in Go, 77
  - combined with replay, 104
  - combining with habits, 77
  - computational difficulty, 26
  - definition, 24
  - hierarchical, 85
  - in animals, 51
  - in rats, 106
- plasticity, 35
  - increasing, 181
- Plato, 160, 175
- pleasure, 53
- policy, 49
- prediction, 55, 56
  - of danger, 92
- prior, 113, 117
  - Bayesian, 114
  - learning, 115
  - subjective, 120
- prioritized sweeping, 105
- Q-value, *see* action-value
- qualia, 137
- rationality, 98
- reinforcement, *see* reward
- reinforcement learning, 48, 104
  - hierarchical, 85
  - options, 85
- relaxation, 193
- relief, 86
- replay
  - active, 177

- and wandering thoughts, 106
  - definition, 104
  - in rats, 106
  - semantic information, 75
  - unsupervised, 105
- resting-state, 100
- reward
  - definition, 47
  - internal, 65
  - perceived, 52, 121, 161
- reward loss
  - attention paid, 158, 187, 191
  - definition, 17, 52
  - perceived, 121, 158, 166
- reward prediction error (RPE), *see* error, reward
  - prediction
- risk aversion, 69
- robot
  - cleaning, 23, 48, 49, 59, 60, 105
  - delivery, 49
  - fetching orange juice, 77
- rumination, 108, 145, 188
- Sāntideva, 199
- Sāriputta, 194
- segmentation, 119
- self, 124
  - and frustration, 158
  - as category, 83, 171
  - as control, 125
  - as desire, 65, 68
  - as seen by other agents, 62
  - based on internal reward, 67
  - destruction, 62
  - esteem and evaluation, 62, 169
  - in Buddhism, 68
  - lack of, *see* no-self
  - needs, 61, 68, 85, 153
    - reducing, 169
  - preservation, 63, 170
  - reducing, 170
  - vs consciousness, 144
  - self-needs, *see* self, needs
  - sensory deprivation, 181
  - shaping, 48
  - signal detection theory, 97
  - simulation, 102, 137, 142, 153, 190
    - reducing, 170, 186
    - taken for real, 143
    - why conscious, 143
- Skeptics, 83, 122, 193
- skills
  - automatization, *see* automatization
  - learning, *see* learning, skills
- Skinner, 134
- social comparison, 62, 170
- social interaction, 199
  - and evaluation, 62
  - as input data, 135
  - basis for consciousness, 138
  - emotions, moral, 95
- social media, 91, 170
- society of mind metaphor, 130
- somatic marker hypothesis, 98
- state of world, 23
- state-value, 49
  - as prediction, 57
  - implications, 57
  - mathematical definition, 49
- stereotypes, 84
- stochastic gradient descent, *see* gradient descent,
  - stochastic
- Stoics, *see* Epictetus
- stress, 18, 69, 163
- subjective experience, 137
  - and pain, 15
  - and suffering, 142, 154
  - in emotions, 90
- subjectivity
  - of perception, 120, 122, 168
  - of thinking, 169
- suffering
  - and consciousness, *see* subjective experience
  - and emotions, 96

- as error signalling, 20, 86
- based on desire in itself, 86
- definition, 17–20
  - as frustration, 17, 20
  - as reward loss, 17
- Buddha, 19
- Cassell, *see* intactness of the person
- Stoics, 19
- van Hooft, 18, 20, 65
- root causes, 149, 163
- suicide, 63, 68, 174
- survival, *see* self, preservation
- symbol, 43
- symbolic AI, *see* GOFAI
- thermostat, 70
- thinking, 102
  - about the future, 102
  - Cassell, 55
  - as planning, 26
  - categorical, 82
  - counterfactual, 52
  - in animals, 106
  - spontaneous, 103
  - with categories, 43
- threat
  - to the intactness of person, *see* intactness of the person
  - definition, 67
- three characteristics (Buddhist), 152, 161
- time scales, 85
- transfer learning, 76
- tree search, 24, 51
  - Monte Carlo, 78
- twelve-link chain, *see* dynamics, Buddhist model
- uncertainty, 69, 151
  - and unpredictability, 166
  - facing it, 163
  - of beliefs, 123
  - of categorization, 82
  - of judgements, 122
  - of perception, 114, 166, 187
  - of reward loss, 121, 122, 166
  - of thoughts, 167, 187
- uncontrollability, 69, 96, 151
  - Buddhism, 125, 162
  - facing it, 161
  - of the mind, 108
  - Stoics, 126
- unexpected behaviour, 46, 57, 64
- unpredictability, 69, 151
  - and uncertainty, 166
  - facing it, 163
- unsatisfactoriness, 152, 164, 165
- unsupervised learning, 41
  - self-supervised, 117
  - with images, 116
- upādāna, 156
- valence
  - Buddhism, 156
  - definition, 93
  - in cognitive dynamics, 154
  - in simulation, 143
- value function, *see* state-value or action-value
- vedanā, 156, 184
- Vedānta, 146
- virtual reality, 137
- vision, 35, 45, 75, 93, 111, *see also* perception
  - as parallel processing, 128
  - difficulty, 111
  - feature extraction, 116
  - illusions, 118
- wandering thoughts, 99
  - and central executive, 130
  - and no-self, 180
  - in meditation, 180, 187
  - increasing suffering, 108, 153
  - observing them, 190
- wanting, *see* desire
- well-being, 160
- wisdom, 199
- Yogacara, 83, 122, 123, 146, 163
  - definition of emptiness, 168
- Zen, 146, 168

# Bibliography

- Abler, B., Walter, H., and Erk, S. (2005). Neural correlates of frustration. *Neuroreport*, 16(7):669–672.
- Ainslie, G. (2001). *Breakdown of will*. Cambridge University Press.
- Aitchison, L. and Lengyel, M. (2016). The Hamiltonian brain: efficient probabilistic inference with excitatory-inhibitory neural circuit dynamics. *PLoS Computational Biology*, 12(12):e1005186.
- Ambrose, R. E., Pfeiffer, B. E., and Foster, D. J. (2016). Reverse replay of hippocampal place cells is uniquely modulated by changing reward. *Neuron*, 91(5):1124–1136.
- Anālayo (2003). *Satipaṭṭhāna: The direct path to realization*. Birmingham: Windhorse Publications.
- Andrews-Hanna, J. R. (2012). The brain’s default network and its adaptive role in internal mentation. *The Neuroscientist*, 18(3):251–270.
- Annas, J. E. and Barnes, J. (1985). *The Modes of Scepticism: Ancient Texts and Modern Interpretations*. Cambridge University Press.
- Arandjelovic, R. and Zisserman, A. (2017). Look, listen and learn. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 609–617. IEEE.
- Arjovsky, M., Bottou, L., Gulrajani, I., and Lopez-Paz, D. (2019). Invariant risk minimization. *arXiv preprint, arXiv:1907.02893*.
- Baars, B. J. (1997). *In the Theater of Consciousness: The Workspace of the Mind*. Oxford University Press.
- Baars, B. J. (2002). The conscious access hypothesis: origins and recent evidence. *Trends in cognitive sciences*, 6(1):47–52.
- Bach, D. R. and Dayan, P. (2017). Algorithms for survival: a comparative perspective on emotions. *Nature Reviews Neuroscience*, 18(5):311–319.
- Baddeley, A. (1996). Exploring the central executive. *The Quarterly Journal of Experimental Psychology Section A*, 49(1):5–28.
- Baer, R. A. (2003). Mindfulness training as a clinical intervention: A conceptual and empirical review. *Clinical psychology: Science and practice*, 10(2):125–143.
- Baird, B., Smallwood, J., and Schooler, J. W. (2011). Back to the future: autobiographical planning and the functionality of mind-wandering. *Consciousness and cognition*, 20(4):1604–1611.
- Balcetis, E. and Dunning, D. (2006). See what you want to see: motivational influences on visual perception. *Journal of personality and social psychology*, 91(4):612.
- Bar, M. (2004). Visual objects in context. *Nature Reviews Neuroscience*, 5(8):617.
- Baumeister, R. F. (1990). Suicide as escape from self. *Psychological review*, 97(1):90.
- Baumeister, R. F., Bratslavsky, E., Finkenauer, C., and Vohs, K. D. (2001). Bad is stronger than good. *Review of general psychology*, 5(4):323.

- Baumeister, R. F., Masicampo, E., and DeWall, C. N. (2009). Prosocial benefits of feeling free: Disbelief in free will increases aggression and reduces helpfulness. *Personality and social psychology bulletin*, 35(2):260–268.
- Baumeister, R. F. and Vohs, K. D. (2002). The pursuit of meaningfulness in life. *Handbook of positive psychology*, 1:608–618.
- Baumeister, R. F., Vohs, K. D., and Tice, D. M. (2007). The strength model of self-control. *Current directions in psychological science*, 16(6):351–355.
- Bavelier, D., Levi, D. M., Li, R. W., Dan, Y., and Hensch, T. K. (2010). Removing brakes on adult brain plasticity: from molecular to behavioral interventions. *Journal of Neuroscience*, 30(45):14964–14971.
- Bechara, A. and Damasio, A. R. (2005). The somatic marker hypothesis: A neural theory of economic decision. *Games and economic behavior*, 52(2):336–372.
- Belk, R. W., Ger, G., and Askegaard, S. (2003). The fire of desire: A multisited inquiry into consumer passion. *Journal of consumer research*, 30(3):326–351.
- Bengio, Y., Ducharme, R., Vincent, P., and Jauvin, C. (2003). A neural probabilistic language model. *Journal of machine learning research*, 3(Feb):1137–1155.
- Beran, M. J., Perner, J., and Proust, J. (2012). *Foundations of metacognition*. Oxford University Press.
- Berkes, P., Orbán, G., Lengyel, M., and Fiser, J. (2011). Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science*, 331(6013):83–87.
- Berridge, K. C. and Kringelbach, M. L. (2015). Pleasure systems in the brain. *Neuron*, 86(3):646–664.
- Beyens, I., Pouwels, J. L., van Driel, I. I., Keijsers, L., and Valkenburg, P. M. (2020). The effect of social media on well-being differs from adolescent to adolescent. *Scientific Reports*, 10(1):1–11.
- Bhatnagar, S., Precup, D., Silver, D., Sutton, R. S., Maei, H. R., and Szepesvári, C. (2009). Convergent temporal-difference learning with arbitrary smooth function approximation. In *Advances in Neural Information Processing Systems*, pages 1204–1212.
- Birch, J. (2016). Natural selection and the maximization of fitness. *Biological Reviews*, 91(3):712–727.
- Birch, J. (2017). Animal sentience and the precautionary principle. *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 2(16):1.
- Bishop, S. R., Lau, M., Shapiro, S., Carlson, L., Anderson, N. D., Carmody, J., Segal, Z. V., Abbey, S., Speca, M., Velting, D., et al. (2004). Mindfulness: a proposed operational definition. *Clinical psychology: Science and practice*, 11(3):230.
- Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and brain sciences*, 18(2):227–247.
- Borji, A., Cheng, M.-M., Jiang, H., and Li, J. (2015). Salient object detection: A benchmark. *IEEE transactions on image processing*, 24(12):5706–5722.
- Bostrom, N. (2003). Are we living in a computer simulation? *The Philosophical Quarterly*, 53(211):243–255.
- Bottou, L. (2003). Stochastic learning. In *Summer School on Machine Learning*, pages 146–168. Springer.
- Botvinick, M. M. (2012). Hierarchical reinforcement learning and decision making. *Current opinion in neurobiology*, 22(6):956–962.
- Botvinick, M. M. and Cohen, J. D. (2014). The computational and neural basis of cognitive control: charted territory and new frontiers. *Cognitive science*, 38(6):1249–1285.
- Bowie, E. (2016). Gymnosophists. In *Oxford Classical Dictionary*. Oxford University Press.
- Brach, T. (2004). *Radical acceptance*. Bantam.



- Brahm, A. (2006). *Mindfulness, Bliss, and Beyond*. Somerville: Wisdom Publications.
- Bratman, M. (1987). *Intention, plans, and practical reason*. CSLI Publications.
- Brewer, J. A., Elwafi, H. M., and Davis, J. H. (2014). Craving to quit: Psychological models and neurobiological mechanisms of mindfulness training as treatment for addictions. *Psychol. Addict. Behav.*, 27(2):366–79.
- Brewer, J. A., Worhunsky, P. D., Gray, J. R., Tang, Y.-Y., Weber, J., and Kober, H. (2011). Meditation experience is associated with differences in default mode network activity and connectivity. *Proceedings of the National Academy of Sciences*, 108(50):20254–20259.
- Brincat, S. L. and Connor, C. E. (2004). Underlying principles of visual shape selectivity in posterior inferotemporal cortex. *Nature neuroscience*, 7(8):880.
- Brochu, E., Cora, V. M., and De Freitas, N. (2010). A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint, arXiv:1012.2599*.
- Brodaric, B. and Neuhaus, F. (2020). Foundations for an ontology of belief, desire and intention. In *Formal Ontology in Information Systems: Proceedings of the 11th International Conference (FOIS 2020)*, volume 330, page 140. IOS Press.
- Brodersen, K. H., Deserno, L., Schlagenhaut, F., Lin, Z., Penny, W. D., Buhmann, J. M., and Stephan, K. E. (2014). Dissecting psychiatric spectrum disorders by generative embedding. *NeuroImage: Clinical*, 4:98–111.
- Brooks, R. A. (1991). Intelligence without representation. *Artificial intelligence*, 47(1-3):139–159.
- Brooks, R. A. (1999). *Cambrian intelligence: The early history of the new AI*. MIT press.
- Brown, N. and Sandholm, T. (2018). Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science*, 359(6374):418–424.
- Browne, C. B., Powley, E., Whitehouse, D., Lucas, S. M., Cowling, P. I., Rohlfshagen, P., Tavener, S., Perez, D., Samothrakis, S., and Colton, S. (2012). A survey of Monte Carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in games*, 4(1):1–43.
- Bruno, M.-A., Vanhaudenhuyse, A., Thibaut, A., Moonen, G., and Laureys, S. (2011). From unresponsive wakefulness to minimally conscious plus and functional locked-in syndromes: recent advances in our understanding of disorders of consciousness. *Journal of neurology*, 258(7):1373–1384.
- Buckner, R. L. (2010). The role of the hippocampus in prediction and imagination. *Annual review of psychology*, 61:27–48.
- Buckner, R. L., Andrews-Hanna, J. R., and Schacter, D. L. (2008). The brain's default network. *Annals of the New York Academy of Sciences*, 1124(1):1–38.
- Bueno-Gómez, N. (2017). Conceptualizing suffering and pain. *Philosophy, Ethics, and Humanities in Medicine*, 12(1):7.
- Bullmore, E. and Sporns, O. (2012). The economy of brain network organization. *Nature Reviews Neuroscience*, 13(5):336.
- Burke, B. L., Martens, A., and Faucher, E. H. (2010). Two decades of terror management theory: A meta-analysis of mortality salience research. *Personality and Social Psychology Review*, 14(2):155–195.
- Butti, C., Santos, M., Uppal, N., and Hof, P. R. (2013). Von Economo neurons: clinical and evolutionary perspectives. *Cortex*, 49(1):312–326.
- Buzsáki, G. (1996). The hippocampo-neocortical dialogue. *Cerebral Cortex*, 6(2):81–92.
- Campbell, M., Hoane Jr, A. J., and Hsu, F.-h. (2002). Deep blue. *Artificial intelligence*, 134(1-2):57–83.

- Carmody, J. (2015). Reconceptualizing mindfulness. *Handbook of mindfulness: Theory, research, and practice*, pages 62–78.
- Carruthers, P. (2009). How we know our own minds: The relationship between mindreading and metacognition. *Behavioral and brain sciences*, 32(2):121–138.
- Carver, C. (2003). Pleasure as a sign you can attend to something else: Placing positive feelings within a general model of affect. *Cognition & Emotion*, 17(2):241–261.
- Cassell, E. (2002). Compassion. In Snyder, C. and Lopez, S., editors, *The Oxford handbook of positive psychology*. Oxford University Press.
- Cassell, E. J. (1982). The nature of suffering and the goals of medicine. *N Engl J Med*, 306:639–645.
- Cassell, E. J. (1989). The relationship between pain and suffering. In Hill, C. S. J. and Fields, W. S., editors, *Advances in Pain Research and Therapy*. Raven Press: New York.
- Castrén, E. and Antila, H. (2017). Neuronal plasticity and neurotrophic factors in drug responses. *Molecular psychiatry*, 22(8):1085.
- Chah, A. (2001). *Being Dharma: The essence of the Buddha's teachings*. Shambhala Publications.
- Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of consciousness studies*, 2(3):200–219.
- Chalmers, D. J. (1996). *The conscious mind: In search of a fundamental theory*. Oxford university press.
- Chan, A. (2018). Laozi. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University.
- Chapman, H. A. and Anderson, A. K. (2012). Understanding disgust. *Annals of the New York Academy of Sciences*, 1251(1):62–76.
- Chaslot, G. M.-B., Winands, M. H., and van Den Herik, H. J. (2008). Parallel Monte Carlo tree search. In *International Conference on Computers and Games*, pages 60–71. Springer.
- Chella, A. and Manzotti, R. (2007). *Artificial consciousness*. Andrews UK Limited.
- Chen, L.-C., Papandreou, G., Kokkinos, I., Murphy, K., and Yuille, A. L. (2018). DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848.
- Chin, J. and Schooler, J. W. (2010). Meta-awareness. *Encyclopedia of consciousness*, 2:33–41.
- Choudhury, S. and Blakemore, S.-J. (2006). Intentions, actions, and the self. In Pockett, S., Banks, W. P., and Gallagher, S., editors, *Does consciousness cause behavior*, pages 39–51. Boston: MIT Press.
- Christoff, K., Gordon, A. M., Smallwood, J., Smith, R., and Schooler, J. W. (2009). Experience sampling during fMRI reveals default network and executive system contributions to mind wandering. *Proceedings of the National Academy of Sciences*, 106(21):8719–8724.
- Clark, A. (2013). Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and brain sciences*, 36(3):181–204.
- Cleeremans, A. and Tallon-Baudry, C. (2021). The function of consciousness is to generate experience. *PsyArXiv*.
- Cohen, M., Quintner, J., and van Rysewyk, S. (2018). Reconsidering the International Association for the Study of Pain definition of pain. *Pain Reports*, 3(2):e634.
- Cohen, P. R. and Levesque, H. J. (1990). Intention is choice with commitment. *Artificial intelligence*, 42(2-3):213–261.

- Colombetti, G. (2005). Appraising valence. *Journal of consciousness studies*, 12(8-9):103–126.
- Colton, S., de Mántaras, R. L., and Stock, O. (2009). Computational creativity: Coming of age. *AI Magazine*, 30(3):11.
- Corballis, M. C. (2019). Language, memory, and mental time travel: An evolutionary perspective. *Frontiers in human neuroscience*, 13:217.
- Corns, J. (2016). Pain eliminativism: scientific and traditional. *Synthese*, 193(9):2949–2971.
- Costa, V. D. and Averbeck, B. B. (2020). Primate orbitofrontal cortex codes information relevant for managing explore–exploit tradeoffs. *Journal of Neuroscience*, 40(12):2553–2561.
- Craig, A. D. (2003). A new view of pain as a homeostatic emotion. *Trends in neurosciences*, 26(6):303–307.
- Craig, A. D. (2009). How do you feel–now? the anterior insula and human awareness. *Nature reviews neuroscience*, 10(1):59–70.
- Crick, F. and Koch, C. (2003). A framework for consciousness. *Nature neuroscience*, 6(2):119.
- Crisp, R. (2017). Well-being. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University.
- Cristia, A., Dupoux, E., Gurven, M., and Stieglitz, J. (2019). Child-directed speech is infrequent in a forager–farmer population: A time allocation study. *Child development*, 90(3):759–773.
- Critchley, H. and Seth, A. (2012). Will studies of macaque insula reveal the neural mechanisms of self-awareness? *Neuron*, 74(3):423–426.
- Csikszentmihalyi, M. (1997). *Finding flow: The psychology of engagement with everyday life*. Basic Books.
- Dahl, C. J., Lutz, A., and Davidson, R. J. (2015). Reconstructing and deconstructing the self: cognitive mechanisms in meditation practice. *Trends in cognitive sciences*, 19(9):515–523.
- Dalai Lama, Lhündrub, K. P., and Cabezón, J. I. (2011). *Meditation on the Nature of Mind*. Boston: Wisdom Publications.
- Damasio, A. (1994). *Descartes’ error: Emotion, Reason, and the Human Brain*. Putnam Publishing.
- Damoiseaux, J., Rombouts, S., Barkhof, F., Scheltens, P., Stam, C., Smith, S. M., and Beckmann, C. (2006). Consistent resting-state networks across healthy subjects. *Proceedings of the national academy of sciences*, 103(37):13848–13853.
- Dasgupta, P. and Maskin, E. (2005). Uncertainty and hyperbolic discounting. *American Economic Review*, 95(4):1290–1299.
- Davidson, J., Liebold, B., Liu, J., Nandy, P., Van Vleet, T., Gargi, U., Gupta, S., He, Y., Lambert, M., Livingston, B., et al. (2010). The YouTube video recommendation system. In *Proceedings of the fourth ACM conference on Recommender systems*, pages 293–296. ACM.
- d’Avila Garcez, A. S., Broda, K. B., and Gabbay, D. M. (2012). *Neural-symbolic learning systems: foundations and applications*. Springer Science & Business Media.
- Daw, N. D., Niv, Y., and Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature neuroscience*, 8(12):1704–1711.
- Dawkins, R. (1986). *The blind watchmaker: Why the evidence of evolution reveals a universe without design*. WW Norton & Company.
- De Berker, A. O., Rutledge, R. B., Mathys, C., Marshall, L., Cross, G. F., Dolan, R. J., and Bestmann, S. (2016). Computations of uncertainty mediate acute stress responses in humans. *Nature Communications*, 7:10996.
- De Brigard, F. and Prinz, J. (2010). Attention and consciousness. *Wiley Interdisciplinary Reviews: Cognitive*

- Science*, 1(1):51–59.
- de Catanzaro, D. (1991). Evolutionary limits to self-preservation. *Ethology and Sociobiology*, 12(1):13–28.
- Degrazia, D. (1998). Suffering. In *Routledge Encyclopedia of Philosophy*. Taylor and Francis.
- Dehaene, S. and Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition*, 79(1-2):1–37.
- Dennett, D. C. (1992). The self as the center of narrative gravity. In Kessel, F., Cole, P., and Johnson, D., editors, *Self and consciousness: Multiple Perspectives*. Hillsdale, NJ: Erlbaum.
- Dhawale, A. K., Smith, M. A., and Ölveczky, B. P. (2017). The role of variability in motor learning. *Annual review of neuroscience*, 40:479–498.
- Diamond, A. (2013). Executive functions. *Annual review of psychology*, 64:135–168.
- Diba, K. and Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nature neuroscience*, 10(10):1241.
- Dietterich, T. G. (2000). Hierarchical reinforcement learning with the MAXQ value function decomposition. *Journal of Artificial Intelligence Research*, 13:227–303.
- Dignum, F., Morley, D., Sonenberg, E. A., and Cavedon, L. (2000). Towards socially sophisticated bdi agents. In *Proceedings fourth international conference on multiagent systems*, pages 111–118. IEEE.
- Diuk, C., Cohen, A., and Littman, M. L. (2008). An object-oriented representation for efficient reinforcement learning. In *Proceedings of the 25th international conference on machine learning (ICML)*, pages 240–247.
- Dolan, R. J. and Dayan, P. (2013). Goals and habits in the brain. *Neuron*, 80(2):312–325.
- Doya, K. and Uchibe, E. (2005). The cyber rodent project: Exploration of adaptive mechanisms for self-preservation and self-reproduction. *Adaptive Behavior*, 13(2):149–160.
- Doyon, J., Penhune, V., and Ungerleider, L. G. (2003). Distinct contribution of the cortico-striatal and cortico-cerebellar systems to motor skill learning. *Neuropsychologia*, 41(3):252–262.
- Drysdale, A. T., Grosenick, L., Downar, J., Dunlop, K., Mansouri, F., Meng, Y., Fetcho, R. N., Zebley, B., Oathes, D. J., Etkin, A., et al. (2017). Resting-state connectivity biomarkers define neurophysiological subtypes of depression. *Nature medicine*, 23(1):28.
- Duffy, K. R. and Mitchell, D. E. (2013). Darkness alters maturation of visual cortex and promotes fast recovery from monocular deprivation. *Current biology*, 23(5):382–386.
- Dunne, J. D., Thompson, E., and Schooler, J. (2019). Mindful meta-awareness: Sustained and non-propositional. *Current Opinion in Psychology*, 28:307–311.
- Edwards, S. D. (2003). Three concepts of suffering. *Medicine, Health Care and Philosophy*, 6(1):59–66.
- Eickenberg, M., Gramfort, A., Varoquaux, G., and Thirion, B. (2017). Seeing it all: Convolutional network layers map the function of the human visual system. *NeuroImage*, 152:184–194.
- Eisenberger, N. I. and Lieberman, M. D. (2004). Why rejection hurts: a common neural alarm system for physical and social pain. *Trends in cognitive sciences*, 8(7):294–300.
- Eisenreich, B. R., Akaishi, R., and Hayden, B. Y. (2017). Control without controllers: toward a distributed neuroscience of executive control. *Journal of cognitive neuroscience*, 29(10):1684–1698.
- Emmons, R. and Shelton, C. (2002). Gratitude and the science of positive psychology. In Snyder, C. and Lopez, S., editors, *The Oxford handbook of positive psychology*. Oxford University Press.
- Evans, J. S. B. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annu. Rev. Psychol.*, 59:255–278.

- Federman, A. (2010). What kind of free will did the Buddha teach? *Philosophy East and West*, pages 1–19.
- Feinman, R. and Lake, B. M. (2018). Learning inductive biases with simple neural networks. *arXiv preprint, arXiv:1802.02745*.
- Feldman, G., Greeson, J., and Senville, J. (2010). Differential effects of mindful breathing, progressive muscle relaxation, and loving-kindness meditation on decentering and negative reactions to repetitive thoughts. *Behaviour research and therapy*, 48(10):1002–1011.
- Fink, G. (2017). Stress: Concepts, definition and history. In *Reference Module in Neuroscience and Biobehavioral Psychology*, pages 1–9. Elsevier.
- Firestone, C. and Scholl, B. J. (2014). "Top-down" effects where none should be found: The El Greco fallacy in perception research. *Psychological science*, 25(1):38–46.
- Fleming, S. M., Dolan, R. J., and Frith, C. D. (2012). Metacognition: computation, biology and function.
- Fletcher, G., editor (2015). *The Routledge handbook of philosophy of well-being*. Routledge.
- Fox, K. C. and Beaty, R. E. (2019). Mind-wandering as creative thinking: neural, psychological, and theoretical considerations. *Current Opinion in Behavioral Sciences*, 27:123–130.
- Fox, K. C., Nijeboer, S., Solomonova, E., Domhoff, G. W., and Christoff, K. (2013). Dreaming as mind wandering: evidence from functional neuroimaging and first-person content reports. *Frontiers in human neuroscience*, 7.
- Fox, M. D. and Raichle, M. E. (2007). Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature reviews neuroscience*, 8(9):700.
- Fraday, E. P., Kent, S. J., Olshausen, B. A., and Sommer, F. T. (2020). Resonator networks, 1: An efficient solution for factoring high-dimensional, distributed representations of data structures. *Neural computation*, 32(12):2311–2331.
- Fredrickson, B. L. (2001). The role of positive emotions in positive psychology: The broaden-and-build theory of positive emotions. *American psychologist*, 56(3):218.
- Freeman, J. B. and Johnson, K. L. (2016). More than meets the eye: Split-second social perception. *Trends in cognitive sciences*, 20(5):362–374.
- Fresco, D. M., Moore, M. T., van Dulmen, M. H., Segal, Z. V., Ma, S. H., Teasdale, J. D., and Williams, J. M. G. (2007). Initial psychometric properties of the experiences questionnaire: validation of a self-report measure of decentering. *Behavior therapy*, 38(3):234–246.
- Friese, M., Messner, C., and Schaffner, Y. (2012). Mindfulness meditation counteracts self-control depletion. *Consciousness and cognition*, 21(2):1016–1022.
- Frijda, N. H. (2016). The evolutionary emergence of what we call "emotions". *Cognition and Emotion*, 30(4):609–620.
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature reviews neuroscience*, 11(2):127.
- Frith, C. (2002). Attention to action and awareness of other minds. *Consciousness and cognition*, 11(4):481–487.
- Frith, C. (2010). What is consciousness for? *Pragmatics & Cognition*, 18(3):497–551.
- Frith, C. (2012). Explaining delusions of control: The comparator model 20 years on. *Consciousness and cognition*, 21(1):52–54.
- Gallagher, S. (2000). Philosophical conceptions of the self: implications for cognitive science. *Trends in cognitive sciences*, 4(1):14–21.
- Gallagher, S. (2003). Phenomenology and experimental design toward a phenomenologically enlightened ex-

- perimental science. *Journal of consciousness studies*, 10(9-10):85–99.
- Gallagher, S. and Shear, J., editors (1999). *Models of the Self*. Imprint Academic.
- Gao, T., Zhou, Y., Li, W., Pfabigan, D. M., and Han, S. (2020). Neural mechanisms of reinforcement learning under mortality threat. *Social neuroscience*, 15(2):170–185.
- Gärdenfors, P. (2004). *Conceptual spaces: The geometry of thought*. MIT press.
- Garfield, J. L. (1990). Epoche and śūnyatā: Skepticism east and west. *Philosophy East and West*, pages 285–307.
- Garfield, J. L. (2010). What is it like to be a bodhisattva? moral phenomenology in śāntideva's bodhicaryāvatāra. *Journal of the International Association of Buddhist Studies*, pages 333–357.
- Garland, E., Froeliger, B., and Howard, M. (2014). Mindfulness training targets neurocognitive mechanisms of addiction at the attention-appraisal-emotion interface. *Frontiers in psychiatry*, 4:173.
- Georgievski, I. and Aiello, M. (2015). HTN planning: Overview, comparison, and beyond. *Artificial Intelligence*, 222:124–156.
- Gershman, S. J. (2017). Reinforcement learning and causal models. *The Oxford handbook of causal reasoning*, page 295.
- Gershman, S. J., Horvitz, E. J., and Tenenbaum, J. B. (2015). Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, 349(6245):273–278.
- Gerstner, W., Lehmann, M., Liakoni, V., Corneil, D., and Brea, J. (2018). Eligibility traces and plasticity on behavioral time scales: experimental support of neohebbian three-factor learning rules. *Frontiers in neural circuits*, 12.
- Gigerenzer, G. and Gaissmaier, W. (2011). Heuristic decision making. *Annual review of psychology*, 62:451–482.
- Goertzel, B. (2012). Perception processing for general intelligence: Bridging the symbolic/subsymbolic gap. In *International Conference on Artificial General Intelligence*, pages 79–88. Springer.
- Goodman, B. and Flaxman, S. (2017). European union regulations on algorithmic decision-making and a "right to explanation". *AI magazine*, 38(3):50–57.
- Gopnik, A. (2009). Could David Hume have known about Buddhism?: Charles François Dolu, the Royal College of La Flèche, and the global Jesuit intellectual network. *Hume Studies*, 35(1/2):5–28.
- Gorder, P. F. (2007). Multicore processors for science and engineering. *Computing in science & engineering*, 9(2):3–7.
- Grabovac, A. D., Lau, M. A., and Willett, B. R. (2011). Mechanisms of mindfulness: A Buddhist psychological model. *Mindfulness*, 2(3):154–166.
- Grafen, A. (2008). The simplest formal argument for fitness optimization. *Journal of genetics*, 87(4):421–433.
- Graser, J. and Stangier, U. (2018). Compassion and loving-kindness meditation: an overview and prospects for the application in clinical samples. *Harvard Review of Psychiatry*, 26(4):201–215.
- Graves, A., Wayne, G., Reynolds, M., Harley, T., Danihelka, I., Grabska-Barwińska, A., Colmenarejo, S. G., Grefenstette, E., Ramalho, T., Agapiou, J., et al. (2016). Hybrid computing using a neural network with dynamic external memory. *Nature*, 538(7626):471.
- Green, D. M. and Swets, J. A. (1988). *Signal detection theory and psychophysics*. Wiley: New York.
- Greff, K., Rasmus, A., Berglund, M., Hao, T., Valpola, H., and Schmidhuber, J. (2016). Tagger: Deep unsupervised perceptual grouping. In *Advances in Neural Information Processing Systems*, pages 4484–4492.
- Gross, C. T. and Canteras, N. S. (2012). The many paths to fear. *Nature Reviews Neuroscience*, 13(9):651.
- Gruber, M. J., Ritchey, M., Wang, S.-F., Doss, M. K., and Ranganath, C. (2016). Post-learning hippocampal

- dynamics promote preferential retention of rewarding events. *Neuron*, 89(5):1110–1120.
- Güçlü, U. and van Gerven, M. A. (2015). Deep neural networks reveal a gradient in the complexity of neural representations across the ventral stream. *Journal of Neuroscience*, 35(27):10005–10014.
- Guestrin, C., Koller, D., Gearhart, C., and Kanodia, N. (2003). Generalizing plans to new environments in relational MDPs. In *Proceedings of the 18th international joint conference on Artificial intelligence*, pages 1003–1010. Morgan Kaufmann Publishers Inc.
- Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., and Pedreschi, D. (2018). A survey of methods for explaining black box models. *ACM Computing Surveys (CSUR)*, 51(5):93.
- Gunaratana, B. H. (2010). *Mindfulness in plain English*. ReadHowYouWant.com.
- Gutmann, M. U., Corander, J., et al. (2016). Bayesian optimization for likelihood-free inference of simulator-based statistical models. *Journal of Machine Learning Research*, 17:1–47.
- Gutmann, M. U., Dutta, R., Kaski, S., and Corander, J. (2018). Likelihood-free inference via classification. *Statistics and Computing*, 28(2):411–425.
- Gutmann, M. U. and Hyvärinen, A. (2012). Noise-contrastive estimation of unnormalized statistical models, with applications to natural image statistics. *J. of Machine Learning Research*, 13:307–361.
- Guyer, P. and Horstmann, R.-P. (2018). Idealism. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University.
- Hadany, L., Beker, T., Eshel, I., and Feldman, M. W. (2006). Why is stress so deadly? An evolutionary perspective. *Proceedings of the Royal Society of London B: Biological Sciences*, 273(1588):881–885.
- Hadot, P. (2002). *Exercices spirituels et philosophie antique*. Albin Michel, revised and augmented edition. Earlier edition partly translated as *Philosophy as a Way of Life*, Wiley, 1995.
- Haggard, P. and Chambon, V. (2012). Sense of agency. *Current Biology*, 22(10):R390–R392.
- Haidt, J. (2003). The moral emotions. In Davidson, R. J., Scherer, K. R., and Goldsmith, H. H., editors, *Handbook of affective sciences*, pages 852–870. Oxford: Oxford University Press.
- Han, C., O’Tuathaigh, C. M., van Trigt, L., Quinn, J. J., Fanselow, M. S., Mongeau, R., Koch, C., and Anderson, D. J. (2003). Trace but not delay fear conditioning requires attention and the anterior cingulate cortex. *Proceedings of the National Academy of Sciences*, 100(22):13087–13092.
- Hardt, M., Recht, B., and Singer, Y. (2016). Train faster, generalize better: Stability of stochastic gradient descent. In *International Conference on Machine Learning*, pages 1225–1234. PMLR.
- Hari, R. (2017). From brain–environment connections to temporal dynamics and social interaction: principles of human brain function. *Neuron*, 94(5):1033–1039.
- Hari, R., Henriksson, L., Malinen, S., and Parkkonen, L. (2015). Centrality of social interaction in human brain function. *Neuron*, 88(1):181–193.
- Hari, R., Parkkonen, L., and Nangini, C. (2010). The brain in time: insights from neuromagnetic recordings. *Annals of the New York Academy of Sciences*, 1191(1):89–109.
- Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena*, 42(1-3):335–346.
- Harnad, S. (2017). To cognize is to categorize: Cognition is categorization. In *Handbook of Categorization in Cognitive Science (Second Edition)*, pages 21–54. Elsevier.
- Harvey, P. (1995). *The selfless mind: Personality, consciousness and nirvana in early Buddhism*. Routledge.
- Harvey, P. (2009). Theravada philosophy of mind and the person. *Buddhist philosophy: Essential readings*, pages 265–74.

- Hassabis, D., Kumaran, D., Vann, S. D., and Maguire, E. A. (2007). Patients with hippocampal amnesia cannot imagine new experiences. *Proceedings of the National Academy of Sciences*, 104(5):1726–1731.
- Hayes, S. C. and Pierson, H. (2005). Acceptance and commitment therapy. In Freeman, A., editor, *Encyclopedia of Cognitive Behavior Therapy*. Springer.
- Hazan, E., Kakade, S., Singh, K., and Van Soest, A. (2019). Provably efficient maximum entropy exploration. In *International Conference on Machine Learning*, pages 2681–2691. PMLR.
- He, H.-Y., Hodos, W., and Quinlan, E. M. (2006). Visual deprivation reactivates rapid ocular dominance plasticity in adult visual cortex. *Journal of Neuroscience*, 26(11):2951–2955.
- Heatherton, T. F. (2011). Neuroscience of self and self-regulation. *Annual review of psychology*, 62:363–390.
- Heatherton, T. F., Wyland, C. L., and Lopez, S. (2003). Assessing self-esteem. *Positive psychological assessment: A handbook of models and measures*, pages 219–233.
- Heathwood, C. (2015). Desire-fulfillment theory. In Fletcher, G., editor, *The Routledge handbook of philosophy of well-being*. Routledge.
- Hebb, D. O. (1949). The organization of behavior; a neuropsychological theory. *A Wiley Book in Clinical Psychology*, pages 62–78.
- Herman, A. (1979). A solution to the paradox of desire in buddhism. *Philosophy East and West*, 29(1):91–94.
- Hesslow, G. (2002). Conscious thought as simulation of behaviour and perception. *Trends in cognitive sciences*, 6(6):242–247.
- Hirsh, J. B., Mar, R. A., and Peterson, J. B. (2012). Psychological entropy: a framework for understanding uncertainty-related anxiety. *Psychological review*, 119(2):304.
- Hirshleifer, J. (1987). On the emotions as guarantors of threats and promises. In *The latest on the best: Essays on evolution and optimality*, pages 307–326. MIT Press: Cambridge, MA.
- Hobfoll, S. E. (1989). Conservation of resources: a new attempt at conceptualizing stress. *American psychologist*, 44(3):513.
- Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Hoffmister, B. (2014). Understanding suffering. In R. Green, N. P., editor, *Suffering and Bioethics*, pages 3–53. Oxford University Press.
- Hofmann, S. G., Grossman, P., and Hinton, D. E. (2011). Loving-kindness and compassion meditation: Potential for psychological interventions. *Clinical psychology review*, 31(7):1126–1132.
- Hofmann, W., Friese, M., and Strack, F. (2009). Impulse and self-control from a dual-systems perspective. *Perspectives on psychological science*, 4(2):162–176.
- Holroyd, C. B. and Coles, M. G. (2002). The neural basis of human error processing: reinforcement learning, dopamine, and the error-related negativity. *Psychological review*, 109(4):679.
- Hölzel, B. K., Lazar, S. W., Gard, T., Schuman-Olivier, Z., Vago, D. R., and Ott, U. (2011). How does mindfulness meditation work? proposing mechanisms of action from a conceptual and neural perspective. *Perspectives on psychological science*, 6(6):537–559.
- Hommel, B. (2013). Dancing in the dark: No role for consciousness in action control. *Frontiers in Psychology*, 4:380.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558.
- Hoyer, P. O. and Hyvärinen, A. (2003). Interpreting neural response variability as Monte Carlo sampling of the



- posterior. In *Advances in neural information processing systems*, pages 293–300.
- Hüllermeier, E. and Waegeman, W. (2021). Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning*, 110(3):457–506.
- Hume, D. (1739). *A Treatise of Human Nature*. Reprinted by Project Gutenberg.
- Huys, Q. J. and Dayan, P. (2009). A Bayesian formulation of behavioral control. *Cognition*, 113(3):314–328.
- Hyvärinen, A., Hoyer, P. O., and Hurri, J. (2009). *Natural Image Statistics: A probabilistic approach to early computational vision*. Springer.
- Hyvärinen, A., Karhunen, J., and Oja, E. (2001). *Independent Component Analysis*. Wiley Interscience.
- Hyvärinen, A. and Morioka, H. (2016). Unsupervised feature extraction by time-contrastive learning and non-linear ICA. In *Advances in Neural Information Processing Systems (NIPS2016)*.
- Hyvärinen, A. and Morioka, H. (2017). Nonlinear ICA of temporally dependent stationary sources. In *Proc. Artificial Intelligence and Statistics (AISTATS2017)*, Fort Lauderdale, Florida.
- Hyvärinen, A. and Oja, E. (1998). Independent component analysis by general nonlinear hebbian-like learning rules. *signal processing*, 64(3):301–313.
- Iacoboni, M. (2005). Neural mechanisms of imitation. *Current opinion in neurobiology*, 15(6):632–637.
- Iannetti, G. D. and Mouraux, A. (2010). From the neuromatrix to the pain matrix (and back). *Experimental brain research*, 205(1):1–12.
- Iannetti, G. D., Salomons, T. V., Moayed, M., Mouraux, A., and Davis, K. D. (2013). Beyond metaphor: contrasting mechanisms of social and physical pain. *Trends in cognitive sciences*, 17(8):371–378.
- Insel, T. R. and Cuthbert, B. N. (2015). Brain disorders? Precisely. *Science*, 348(6234):499–500.
- Inzlicht, M., Shenav, A., and Olivola, C. Y. (2018). The effort paradox: Effort is both costly and valued. *Trends in cognitive sciences*, 22(4):337–349.
- Jacob, J., Kent, M., Benson-Amram, S., Herculano-Houzel, S., Raghanti, M. A., Ploppert, E., Drake, J., Hindi, B., Natale, N. R., Daniels, S., et al. (2021). Cytoarchitectural characteristics associated with cognitive flexibility in raccoons. *Journal of Comparative Neurology*.
- Janet, P. (1889). *L'Automatisme psychologique*. Paris: Felix Alcan.
- Kabat-Zinn, J. (2012). *Mindfulness for beginners: Reclaiming the present moment—and your life*. Sounds True.
- Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux: New York.
- Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., et al. (2019). Advances and open problems in federated learning. *arXiv preprint, arXiv:1912.04977*.
- Kane, M. J., Brown, L. H., McVay, J. C., Silvia, P. J., Myin-Germeys, I., and Kwapil, T. R. (2007). For whom the mind wanders, and when: An experience-sampling study of working memory and executive control in daily life. *Psychological science*, 18(7):614–621.
- Karlsson, M. P. and Frank, L. M. (2009). Awake replay of remote experiences in the hippocampus. *Nature Neuroscience*, 12(7):913–918.
- Kasser, T. and Sheldon, K. M. (2000). Of wealth and death: Materialism, mortality salience, and consumption behavior. *Psychological science*, 11(4):348–351.
- Kavanagh, D. J., Andrade, J., and May, J. (2005). Imaginary relish and exquisite torture: the elaborated intrusion theory of desire. *Psychological review*, 112(2):446.
- Keramati, M. and Gutkin, B. (2014). Homeostatic reinforcement learning for integrating reward collection and

- physiological stability. *Elife*, 3:e04811.
- Kersten, D., Mamassian, P., and Yuille, A. (2004). Object perception as Bayesian inference. *Annual Review of Psychology*, 55:271–304.
- Khemakhem, I., Kingma, D. P., Monti, R. P., and Hyvärinen, A. (2020). Variational autoencoders and nonlinear ICA: A unifying framework. In *Proc. Artificial Intelligence and Statistics (AISTATS2020)*.
- Killingsworth, M. A. and Gilbert, D. T. (2010). A wandering mind is an unhappy mind. *Science*, 330(6006):932–932.
- Kirchner, H. and Thorpe, S. J. (2006). Ultra-rapid object detection with saccadic eye movements: Visual processing speed revisited. *Vision research*, 46(11):1762–1776.
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526.
- Kiviniemi, V., Kantola, J.-H., Jauhiainen, J., Hyvärinen, A., and Tervonen, O. (2003). Independent component analysis of nondeterministic fMRI signal sources. *Neuroimage*, 19(2):253–260.
- Klein, C. (2007). An imperative theory of pain. *The Journal of Philosophy*, 104(10):517–532.
- Klinger, E. (2013). Goal commitments and the content of thoughts and dreams: Basic principles. *Frontiers in Psychology*, 4:415.
- Koechlin, E. and Summerfield, C. (2007). An information theoretical approach to prefrontal executive function. *Trends in cognitive sciences*, 11(6):229–235.
- Konstan, D. (2018). Epicurus. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University.
- Koolhaas, J., Bartolomucci, A., Buwalda, B. d., De Boer, S., Flügge, G., Korte, S., Meerlo, P., Murison, R., Olivier, B., Palanza, P., et al. (2011). Stress revisited: a critical evaluation of the stress concept. *Neuroscience & Biobehavioral Reviews*, 35(5):1291–1301.
- Kording, K. P., Tenenbaum, J. B., and Shadmehr, R. (2007). The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nature neuroscience*, 10(6):779.
- Korenberg, A. T. and Ghahramani, Z. (2002). A Bayesian view of motor adaptation. *Current Psychology of Cognition*, 21(4/5):537–564.
- Kőszegi, B. and Rabin, M. (2006). A model of reference-dependent preferences. *The Quarterly Journal of Economics*, 121(4):1133–1165.
- Kouider, S., De Gardelle, V., Sackur, J., and Dupoux, E. (2010). How rich is consciousness? The partial awareness hypothesis. *Trends in cognitive sciences*, 14(7):301–307.
- Kozachkov, L., Lundqvist, M., Slotine, J.-J., and Miller, E. K. (2020). Achieving stable dynamics in neural circuits. *PLoS computational biology*, 16(8):e1007659.
- Kriegeskorte, N. (2015). Deep neural networks: a new framework for modeling biological vision and brain information processing. *Annual review of vision science*, 1:417–446.
- Krueger, K. A. and Dayan, P. (2009). Flexible shaping: How learning in small steps helps. *Cognition*, 110(3):380–394.
- Kruger, J., Wirtz, D., Van Boven, L., and Altermatt, T. W. (2004). The effort heuristic. *Journal of Experimental Social Psychology*, 40(1):91–98.
- Kurth-Nelson, Z., Economides, M., Dolan, R. J., and Dayan, P. (2016). Fast sequences of non-spatial state rep-

- representations in humans. *Neuron*, 91(1):194–204.
- Kuyken, W., Watkins, E., Holden, E., White, K., Taylor, R. S., Byford, S., Evans, A., Radford, S., Teasdale, J. D., and Dalgleish, T. (2010). How does mindfulness-based cognitive therapy work? *Behaviour research and therapy*, 48(11):1105–1112.
- Kwapis, J. L., Jarome, T. J., and Helmstetter, F. J. (2015). The role of the medial prefrontal cortex in trace fear extinction. *Learning & Memory*, 22(1):39–46.
- Kyabgon, T. (2015). *Moonbeams of Mahamudra*. Carlton North: Shogam Publications.
- Lake, B. M., Salakhutdinov, R., and Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338.
- Lambie, G. W. and Haugen, J. S. (2019). Understanding greed as a unified construct. *Personality and Individual Differences*, 141:31–39.
- Lamme, V. A. and Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in neurosciences*, 23(11):571–579.
- Lao, S.-A., Kissane, D., and Meadows, G. (2016). Cognitive effects of mbsr/mbct: A systematic review of neuropsychological outcomes. *Consciousness and cognition*, 45:109–123.
- Larsson, G., Maire, M., and Shakhnarovich, G. (2017). Colorization as a proxy task for visual understanding. In *CVPR*, volume 2, page 8.
- Lau, H. and Rosenthal, D. (2011). Empirical support for higher-order theories of conscious awareness. *Trends in cognitive sciences*, 15(8):365–373.
- Lazarus, R. S. (1993). From psychological stress to the emotions: A history of changing outlooks. *Annual review of psychology*, 44(1):1–22.
- Lazarus, R. S. and Folkman, S. (1984). *Stress, appraisal, and coping*. Springer publishing company.
- LeDoux, J. E. and Pine, D. S. (2016). Using neuroscience to help understand fear and anxiety: a two-system framework. *American Journal of Psychiatry*, 173(11):1083–1093.
- Legenstein, R., Chase, S. M., Schwartz, A. B., and Maass, W. (2010). A reward-modulated hebbian learning rule can explain experimentally observed network reorganization in a brain control task. *Journal of Neuroscience*, 30(25):8400–8410.
- Legg, S. and Hutter, M. (2007). Universal intelligence: A definition of machine intelligence. *Minds and Machines*, 17(4):391–444.
- Leknes, S., Brooks, J. C., Wiech, K., and Tracey, I. (2008). Pain relief as an opponent process: a psychophysical investigation. *European journal of neuroscience*, 28(4):794–801.
- Leknes, S. and Tracey, I. (2008). A common neurobiology for pain and pleasure. *Nature Reviews Neuroscience*, 9(4):314–320.
- Lengyel, M. and Dayan, P. (2008). Hippocampal contributions to control: the third way. In *Advances in neural information processing systems*, pages 889–896.
- Libet, B., Gleason, C. A., Wright, E. W., and Pearl, D. K. (1983). Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential) the unconscious initiation of a freely voluntary act. *Brain*, 106(3):623–642.
- Lieder, F. and Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, 43.
- Lin, L. J. (1991). Programming robots using reinforcement learning and teaching. In *AAAI*, pages 781–786.

- Lindsay, E. K. and Creswell, J. D. (2017). Mechanisms of mindfulness training: Monitor and acceptance theory (MAT). *Clinical Psychology Review*, 51:48–59.
- Liu, Y.-Y., Slotine, J.-J., and Barabási, A.-L. (2011). Controllability of complex networks. *Nature*, 473(7346):167.
- Livneh, Y., Ramesh, R. N., Burgess, C. R., Levandowski, K. M., Madara, J. C., Fenselau, H., Goldey, G. J., Diaz, V. E., Jikomes, N., Resch, J. M., et al. (2017). Homeostatic circuits selectively gate food cue responses in insular cortex. *Nature*, 546(7660):611.
- Lu, H., Zou, Q., Gu, H., Raichle, M. E., Stein, E. A., and Yang, Y. (2012). Rat brains also have a default mode network. *Proceedings of the National Academy of Sciences*, 109(10):3979–3984.
- Lusthaus, D. (1998). Buddhism, Yogācāra school of. In *Routledge Encyclopedia of Philosophy*. Taylor and Francis.
- Lusthaus, D. (2013). What is and isn't Yogācāra. Yogacara Buddhism Research Association.
- Lutz, A., Slagter, H. A., Dunne, J. D., and Davidson, R. J. (2008). Attention regulation and monitoring in meditation. *Trends in cognitive sciences*, 12(4):163–169.
- Lutz, A. and Thompson, E. (2003). Neurophenomenology integrating subjective experience and brain dynamics in the neuroscience of consciousness. *Journal of consciousness studies*, 10(9-10):31–52.
- Ly, C., Greb, A. C., Cameron, L. P., Wong, J. M., Barragan, E. V., Wilson, P. C., Burbach, K. F., Zarandi, S. S., Sood, A., Paddy, M. R., et al. (2018). Psychedelics promote structural and functional neural plasticity. *Cell reports*, 23(11):3170–3182.
- Lyubomirsky, S. (2010). Hedonic adaptation to positive and negative experiences. *The Oxford handbook of stress, health, and coping*.
- Ma, W. J., Kording, K. P., and Goldreich, D. (2022). *Bayesian models of perception and action*. MIT Press. In press.
- MacDonald, G. (2009). Social pain and hurt feelings. *Cambridge handbook of personality psychology*, pages 541–555.
- MacKenzie, M. J. and Baumeister, R. F. (2015). Self-regulatory strength and mindfulness. In *Handbook of mindfulness and self-regulation*, pages 95–105. Springer.
- Mahasi, S. (1996). *Great Discourse on Not Self (Anattalakkhana Sutta)*. Bangkok: Buddhadhamma Foundation.
- Mahasi, S. (1999). *A Discourse on Dependent Origination*. Bangkok: Buddhadhamma Foundation.
- Mahasi, S. (2016). *Manual of Insight*. Simon & Schuster.
- Mai, V., Mani, K., and Paull, L. (2022). Sample efficient deep reinforcement learning via uncertainty estimation. *arXiv preprint arXiv:2201.01666*.
- Malt, B. C., Ross, B. H., and Murphy, G. L. (1995). Predicting features for members of natural categories when categorization is uncertain. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21(3):646.
- Mancinelli, F., Roiser, J., and Dayan, P. (2021). Internality and the internalisation of failure: Evidence from a novel task. *PLoS Computational Biology*, 17(7):e1009134.
- Marchetti, I., Koster, E. H., Klinger, E., and Alloy, L. B. (2016). Spontaneous thought and vulnerability to mood disorders: the dark side of the wandering mind. *Clinical Psychological Science*, 4(5):835–857.
- Marchetti, I., Van de Putte, E., and Koster, E. H. (2014). Self-generated thoughts and depression: from day-dreaming to depressive symptoms. *Frontiers in human neuroscience*, 8:131.
- Marek, S. and Dosenbach, N. U. (2018). The frontoparietal network: function, electrophysiology, and impor-

- tance of individual precision mapping. *Dialogues in clinical neuroscience*, 20(2):133.
- Markov, I. L. (2014). Limits on fundamental limits to computation. *Nature*, 512(7513):147.
- Markram, H., Gerstner, W., and Sjöström, P. J. (2012). Spike-timing-dependent plasticity: a comprehensive overview. *Frontiers in synaptic neuroscience*, 4:2.
- Martela, F. (2020). *A Wonderful Life: Insights on Finding a Meaningful Existence*. HarperCollins.
- Martin, D., Fowlkes, C., Tal, D., and Malik, J. (2001). A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proc. 8th Int'l Conf. Computer Vision*, volume 2, pages 416–423.
- Mates, B. (1996). *The skeptic way: Sextus Empiricus' outlines of Pyrrhonism*. Oxford University Press.
- Mattar, M. G. and Daw, N. D. (2018). Prioritized memory access explains planning and hippocampal replay. *Nature Neuroscience*, 21(11):1609–1617.
- McCloskey, M. and Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier.
- McCullough, M. and vanOyen-Witvliet, C. (2002). The psychology of forgiveness. In Snyder, C. and Lopez, S., editors, *The Oxford handbook of positive psychology*. Oxford University Press.
- McDermott, D. (2007). Artificial intelligence and consciousness. *The Cambridge handbook of consciousness*, pages 117–150.
- McEvilley, T. (1982). Pyrrhonism and mādhyamika. *Philosophy East and West*, pages 3–35.
- Mee, S., Bunney, B. G., Reist, C., Potkin, S. G., and Bunney, W. E. (2006). Psychological pain: a review of evidence. *Journal of Psychiatric Research*, 40(8):680–690.
- Meerwijk, E. L. and Weiss, S. J. (2011). Toward a unifying definition of psychological pain. *Journal of Loss and Trauma*, 16(5):402–412.
- Mendel, J. M. (1995). Fuzzy logic systems for engineering: a tutorial. *Proceedings of the IEEE*, 83(3):345–377.
- Metzinger, T. (2003). *Being no one*. Cambridge, MA: MIT Press.
- Mikolov, T., Karafiát, M., Burget, L., Černocký, J., and Khudanpur, S. (2010). Recurrent neural network based language model. In *Eleventh annual conference of the international speech communication association*.
- Miller, R. R., Barnet, R. C., and Grahame, N. J. (1995). Assessment of the Rescorla–Wagner model. *Psychological bulletin*, 117(3):363.
- Minsky, M. (1988). *Society of mind*. Simon and Schuster.
- Minut, S. and Mahadevan, S. (2001). A reinforcement learning model of selective visual attention. In *Proceedings of the fifth international conference on autonomous agents*, pages 457–464. ACM.
- Mirolli, M. and Baldassarre, G. (2013). Intrinsically motivated learning in natural and artificial systems. *Intrinsically Motivated Learning in Natural and Artificial Systems*, pages 49–72.
- Misra, I., Zitnick, C. L., and Hebert, M. (2016). Shuffle and learn: unsupervised learning using temporal order verification. In *European Conference on Computer Vision*, pages 527–544. Springer.
- Miyake, A., Friedman, N. P., Emerson, M. J., Witzki, A. H., Howerter, A., and Wager, T. D. (2000). The unity and diversity of executive functions and their contributions to complex "frontal lobe" tasks: A latent variable analysis. *Cognitive psychology*, 41(1):49–100.
- Mnih, V., Heess, N., Graves, A., et al. (2014). Recurrent models of visual attention. In *Advances in neural information processing systems*, pages 2204–2212.
- Momennejad, I., Otto, A. R., Daw, N. D., and Norman, K. A. (2018). Offline replay supports planning in human

- reinforcement learning. *eLife*, e32548.
- Monti, M. M., Vanhaudenhuyse, A., Coleman, M. R., Boly, M., Pickard, J. D., Tshibanda, L., Owen, A. M., and Laureys, S. (2010). Willful modulation of brain activity in disorders of consciousness. *New England Journal of Medicine*, 362(7):579–589.
- Moore, A. W. and Atkeson, C. G. (1993). Prioritized sweeping: Reinforcement learning with less data and less time. *Machine learning*, 13(1):103–130.
- Moreno-Bote, R., Knill, D. C., and Pouget, A. (2011). Bayesian sampling in visual perception. *Proceedings of the National Academy of Sciences*, 108(30):12491–12496.
- Morison, B. (2019). Sextus Empiricus. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University.
- Mulder, J. M. (2018). Why intentions? *Ratio*, 31:51–64.
- Murphy, G. L. and Ross, B. H. (2010). Uncertainty in category-based induction: When do people integrate across categories? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36(2):263.
- Nakayama, K. (1999). Mid-level vision. *The MIT encyclopedia of the cognitive sciences*.
- Nau, D. S., Au, T.-C., Ilghami, O., Kuter, U., Murdock, J. W., Wu, D., and Yaman, F. (2003). SHOP2: an HTN planning system. *Journal of artificial intelligence research*, 20:379–404.
- Neff, K. D., Kirkpatrick, K. L., and Rude, S. S. (2007). Self-compassion and adaptive psychological functioning. *Journal of research in personality*, 41(1):139–154.
- Nesse, R. M. (2000). Is depression an adaptation? *Archives of General Psychiatry*, 57(1):14–20.
- Nevin, J. A. (1999). Analyzing Thorndike’s law of effect: The question of stimulus-response bonds. *Journal of the experimental analysis of behavior*, 72(3):447–450.
- Ng, A. Y., Harada, D., and Russell, S. (1999). Policy invariance under reward transformations: Theory and application to reward shaping. In *ICML*, volume 99, pages 278–287.
- NIDA (2020). How does cocaine produce its effects? National Institute on Drug Abuse, retrieved from <https://www.drugabuse.gov/publications/research-reports/cocaine/how-does-cocaine-produce-its-effects>, on 26 Oct 2021.
- Nisargadatta, M. (1982). *Seeds of Consciousness*. New York: Grove Press.
- Niv, Y. (2009). Reinforcement learning in the brain. *Journal of Mathematical Psychology*, 53(3):139–154.
- Nokia, M. S., Lensu, S., Ahtiainen, J. P., Johansson, P. P., Koch, L. G., Britton, S. L., and Kainulainen, H. (2016). Physical exercise increases adult hippocampal neurogenesis in male rats provided it is aerobic and sustained. *The Journal of physiology*, 594(7):1855–1873.
- Nowak, A., Gelfand, M. J., Borkowski, W., Cohen, D., and Hernandez, I. (2016). The evolutionary basis of honor cultures. *Psychological science*, 27(1):12–24.
- Nowak, M. A., Tarnita, C. E., and Wilson, E. O. (2010). The evolution of eusociality. *Nature*, 466(7310):1057.
- Nummenmaa, L., Glerean, E., Hari, R., and Hietanen, J. K. (2014). Bodily maps of emotions. *Proceedings of the National Academy of Sciences*, 111(2):646–651.
- Nummenmaa, L., Glerean, E., Viinikainen, M., Jääskeläinen, I. P., Hari, R., and Sams, M. (2012). Emotions promote social interaction by synchronizing brain activity across individuals. *Proceedings of the National Academy of Sciences*, 109(24):9599–9604.
- Nutt, D. J., Lingford-Hughes, A., Erritzoe, D., and Stokes, P. R. (2015). The dopamine theory of addiction: 40 years of highs and lows. *Nature Reviews Neuroscience*, 16(5):305.

- Oatley, K. and Johnson-Laird, P. N. (1987). Towards a cognitive theory of emotions. *Cognition and emotion*, 1(1):29–50.
- Oatley, K. and Johnson-Laird, P. N. (1990). Semantic primitives for emotions: A reply to Ortony and Clore. *Cognition and Emotion*, 4(2):129–143.
- Oja, E. (1982). Simplified neuron model as a principal component analyzer. *Journal of mathematical biology*, 15(3):267–273.
- Oja, E. (1992). Principal components, minor components, and linear neural networks. *Neural networks*, 5(6):927–935.
- Oldmeadow, H. (1997). Delivering the last blade of grass: Aspects of the bodhisattva ideal in the mahāyāna. *Asian Philosophy*, 7(3):181–194.
- Olshausen, B. A. and Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607.
- OpenAI (2018). OpenAI Five. <https://blog.openai.com/openai-five/>.
- Orbach, I., Mikulincer, M., Sirota, P., and Gilboa-Schechtman, E. (2003). Mental pain: a multidimensional operationalization and definition. *Suicide and Life-Threatening Behavior*, 33(3):219–230.
- O’Reilly, R. C., Hazy, T. E., Mollick, J., Mackie, P., and Herd, S. (2014). Goal-driven cognition in the brain: a computational framework. *arXiv preprint, arXiv:1404.7591*.
- Ottaviani, C., Shapiro, D., and Couyoumdjian, A. (2013). Flexibility as the key for somatic health: From mind wandering to perseverative cognition. *Biological psychology*, 94(1):38–43.
- Palmwood, E. N. and McBride, C. A. (2019). Challenge vs. threat: The effect of appraisal type on resource depletion. *Current Psychology*, 38(6):1522–1529.
- Pan, S. J. and Yang, Q. (2009). A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359.
- Papies, E. K. and Barsalou, L. W. (2015). Grounding desire and motivated behavior: A theoretical framework and review of empirical evidence. *The psychology of desire*, pages 36–60.
- Papini, M. R., Fuchs, P. N., and Torres, C. (2015). Behavioral neuroscience of psychological pain. *Neuroscience & Biobehavioral Reviews*, 48:53–69.
- Parker, G. A. and Smith, J. M. (1990). Optimality theory in evolutionary biology. *Nature*, 348(6296):27.
- Pascanu, R., Li, Y., Vinyals, O., Heess, N., Buesing, L., Racanière, S., Reichert, D., Weber, T., Wierstra, D., and Battaglia, P. (2017). Learning model-based planning from scratch. *arXiv preprint, arXiv:1707.06170*.
- Pascual-Leone, A., Amedi, A., Fregni, F., and Merabet, L. B. (2005). The plastic human brain cortex. *Annu. Rev. Neurosci.*, 28:377–401.
- Pathak, D., Agrawal, P., Efros, A. A., and Darrell, T. (2017). Curiosity-driven exploration by self-supervised prediction. In *International Conference on Machine Learning (ICML)*.
- Peacock, J. (2018). Vedanā, ethics and character: A prolegomena. *Contemporary Buddhism*, 19(1):160–184.
- Pearl, J. (2009). *Causality*. Cambridge University Press.
- Peralta, V. and Cuesta, M. J. (2001). How many and which are the psychopathological dimensions in schizophrenia? Issues influencing their ascertainment. *Schizophrenia research*, 49(3):269–285.
- Perea, G., Navarrete, M., and Araque, A. (2009). Tripartite synapses: astrocytes process and control synaptic information. *Trends in neurosciences*, 32(8):421–431.
- Peters, J., Janzing, D., and Schölkopf, B. (2017). *Elements of causal inference: foundations and learning algo-*

- rhythms*. The MIT Press.
- Peters, J., Mülling, K., Kober, J., Nguyen-Tuong, D., and Krömer, O. (2011). Towards motor skill learning for robotics. In *Robotics Research*, pages 469–482. Springer.
- Peterson, C. (1999). Personal control and well-being. In *Well-being: Foundations of hedonic psychology*. Russell Sage Foundation.
- Pfeiffer, B. E. and Foster, D. J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*, 497(7447):74–79.
- Pinker, S. (1999). How the mind works. *Annals of the New York Academy of Sciences*, 882(1):119–127.
- Pisella, L., Rode, G., Farne, A., Tilikete, C., and Rossetti, Y. (2006). Prism adaptation in the rehabilitation of patients with visuo-spatial cognitive disorders. *Current opinion in neurology*, 19(6):534–542.
- Platt, M. L. and Huettel, S. A. (2008). Risky business: the neuroeconomics of decision making under uncertainty. *Nature neuroscience*, 11(4):398.
- Pockett, S., Banks, W. P., and Gallagher, S., editors (2009). *Does consciousness cause behavior?* MIT Press.
- Poerio, G. L., Totterdell, P., Emerson, L.-M., and Miles, E. (2015). Love is the triumph of the imagination: Day-dreams about significant others are associated with increased happiness, love and connection. *Consciousness and Cognition*, 33:135–144.
- Poerio, G. L., Totterdell, P., and Miles, E. (2013). Mind-wandering and negative mood: Does one thing really lead to another? *Consciousness and cognition*, 22(4):1412–1421.
- Poole, D. L. and Mackworth, A. K. (2010). *Artificial Intelligence: foundations of computational agents*. Cambridge University Press. Free online at <https://artint.info/2e/>.
- Pramote, P. (2013). *A Path to Enlightenment I*. Prima Publishing: Bangkok.
- Prebble, S. C., Addis, D. R., and Tippett, L. J. (2013). Autobiographical memory and sense of self. *Psychological bulletin*, 139(4):815.
- Proust, J. (2010). Metacognition. *Philosophy Compass*, 5(11):989–998.
- Quiroga, R. Q., Kreiman, G., Koch, C., and Fried, I. (2008). Sparse but not 'grandmother-cell' coding in the medial temporal lobe. *Trends in cognitive sciences*, 12(3):87–91.
- Quiroga, R. Q., Reddy, L., Kreiman, G., Koch, C., and Fried, I. (2005). Invariant visual representation by single neurons in the human brain. *Nature*, 435(7045):1102.
- Raffaelli, W. and Arnaudo, E. (2017). Pain as a disease: an overview. *Journal of pain research*, 10:2003.
- Raichle, M. E. (2015). The brain's default mode network. *Annual review of neuroscience*, 38:433–447.
- Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., and Shulman, G. L. (2001). A default mode of brain function. *Proceedings of the National Academy of Sciences*, 98(2):676–682.
- Raja, S. N., Carr, D. B., Cohen, M., Finnerup, N. B., Flor, H., Gibson, S., Keefe, F. J., Mogil, J. S., Ringkamp, M., Sluka, K. A., et al. (2020). The revised international association for the study of pain definition of pain: concepts, challenges, and compromises. *Pain*, 161(9):1976–1982.
- Rao, A. S. and Georgeff, M. P. (1991). Modeling rational agents within a BDI-architecture. *KR*, 91:473–484.
- Recht, B., Re, C., Wright, S., and Niu, F. (2011). Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *Advances in neural information processing systems*, pages 693–701.
- Recht, B., Roelofs, R., Schmidt, L., and Shankar, V. (2019). Do imagenet classifiers generalize to imagenet? In *International Conference on Machine Learning*, pages 5389–5400. PMLR.
- Redgrave, P. and Gurney, K. (2006). The short-latency dopamine signal: a role in discovering novel actions?



- Nature reviews neuroscience*, 7(12):967.
- Redshaw, J. and Bulley, A. (2018). Future-thinking in animals. *The psychology of thinking about the future*, page 31.
- Reggia, J. A. (2013). The rise of machine consciousness: Studying consciousness with computational models. *Neural Networks*, 44:112–131.
- Revonsuo, A. (2000). The reinterpretation of dreams: An evolutionary hypothesis of the function of dreaming. *Behavioral and Brain Sciences*, 23(6):877–901.
- Roese, N. J. (1997). Counterfactual thinking. *Psychological bulletin*, 121(1):133.
- Rosch, E. (1978). Principles of categorization. *Cognition and categorization*.
- Rosch, E. (1999). Reclaiming concepts. *Journal of consciousness studies*, 6(11-12):61–77.
- Rosenberg, A. and Bouchard, F. (2011). Fitness. In Zalta, E. N., editor, *Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University.
- Roskies, A. (2006). Neuroscientific challenges to free will and responsibility. *Trends in cognitive sciences*, 10(9):419–423.
- Ruby, F. J. M., Smallwood, J., Engen, H., and Singer, T. (2013). How self-generated thought shapes mood — the relation between mind-wandering and mood depends on the socio-temporal content of thoughts. *PLoS ONE*, 8(10):1–7.
- Rueda, M. R., Posner, M. I., and Rothbart, M. K. (2004). Attentional control and self-regulation. *Handbook of self-regulation: Research, theory, and applications*, 2:284–299.
- Russell, S. and Norvig, P. (2020). *Artificial intelligence: a modern approach*. Pearson Education, 4th edition.
- Russell, S. J. (1997). Rationality and intelligence. *Artificial intelligence*, 94(1-2):57–77.
- Rutledge, R. B., De Berker, A. O., Espenhahn, S., Dayan, P., and Dolan, R. J. (2016). The social contingency of momentary subjective well-being. *Nature Communications*, 7:11825.
- Rutledge, R. B., Skandali, N., Dayan, P., and Dolan, R. J. (2014). A computational and neural model of momentary subjective well-being. *Proceedings of the National Academy of Sciences*, 111(33):12252–12257.
- Safran, J. and Segal, Z. V. (1990). *Interpersonal process in cognitive therapy*. Jason Aronson, Inc.
- Sagi, D. (2011). Perceptual learning in vision research. *Vision research*, 51(13):1552–1566.
- Sale, A., Vetencourt, J. F. M., Medini, P., Cenni, M. C., Baroncelli, L., De Pasquale, R., and Maffei, L. (2007). Environmental enrichment in adulthood promotes amblyopia recovery through a reduction of intracortical inhibition. *Nature neuroscience*, 10(6):679.
- Salimans, T., Ho, J., Chen, X., Sidor, S., and Sutskever, I. (2017). Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint, arXiv:1703.03864*.
- Salzberg, S. (2002). *Lovingkindness: The revolutionary art of happiness*. Shambhala Publications.
- Samuel, A. (1959). Some studies in machine learning using the game of checkers. *IBM J. of Research and Development*, 3:210–229.
- Sapolsky, R. M. (2004). *Why zebras don't get ulcers*. Holt paperbacks, 2nd edition.
- Schaal, S. (1999). Is imitation learning the route to humanoid robots? *Trends in cognitive sciences*, 3(6):233–242.
- Schaul, T., Quan, J., Antonoglou, I., and Silver, D. (2016). Prioritized experience replay. In *International Conference on Learning Representations*.
- Scherer, K. R. (2009). The dynamic architecture of emotion: Evidence for the component process model. *Cognition and emotion*, 23(7):1307–1351.

- Scherer, K. R. (2011). On the rationality of emotions: or, when are emotions rational? *Social Science Information*, 50(3-4):330–350.
- Schmidhuber, J. (1991). Curious model-building control systems. In *Neural Networks, 1991. 1991 IEEE International Joint Conference on*, pages 1458–1463. IEEE.
- Schooler, J. W., Smallwood, J., Christoff, K., Handy, T. C., Reichle, E. D., and Sayette, M. A. (2011). Meta-awareness, perceptual decoupling and the wandering mind. *Trends in cognitive sciences*, 15(7):319–326.
- Schroeder, J. W. (2004). *Skillful means: The heart of Buddhist compassion*, volume 54. Motilal Banarsidass Publ.
- Schroeder, T. (2017). Desire. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University.
- Schultz, W. (2016). Dopamine reward prediction error coding. *Dialogues in clinical neuroscience*, 18(1):23.
- Schwartz, B. (2004). *The paradox of choice: Why more is less*. New York: HarperCollins.
- Sebastian, C., Burnett, S., and Blakemore, S.-J. (2008). Development of the self-concept during adolescence. *Trends in cognitive sciences*, 12(11):441–446.
- Seth, A. (2009). Functions of consciousness. In Banks, W., editor, *Encyclopedia of Consciousness*, volume 1, pages 279–293. Elsevier Press.
- Seth, A. K., Baars, B. J., and Edelman, D. B. (2005). Criteria for consciousness in humans and other mammals. *Consciousness and cognition*, 14(1):119–139.
- Seth, A. K. and Tsakiris, M. (2018). Being a beast machine: the somatic basis of selfhood. *Trends in cognitive sciences*, 22(11):969–981.
- Seymour, B. (2019). Pain: a precision signal for reinforcement learning and control. *Neuron*, 101(6):1029–1041.
- Seymour, B., O’Doherty, J. P., Koltzenburg, M., Wiech, K., Frackowiak, R., Friston, K., and Dolan, R. (2005). Opponent appetitive-aversive neural processes underlie predictive learning of pain relief. *Nature neuroscience*, 8(9):1234–1240.
- Shapiro, S. L., Carlson, L. E., Astin, J. A., and Freedman, B. (2006). Mechanisms of mindfulness. *Journal of clinical psychology*, 62(3):373–386.
- Shulman, G. L., Fiez, J. A., Corbetta, M., Buckner, R. L., Miezin, F. M., Raichle, M. E., and Petersen, S. E. (1997). Common blood flow changes across visual tasks: II. decreases in cerebral cortex. *Journal of cognitive neuroscience*, 9(5):648–663.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. (2017). Mastering the game of Go without human knowledge. *Nature*, 550(7676):354.
- Simon, H. A. (1967). Motivational and emotional controls of cognition. *Psychological review*, 74(1):29.
- Simon, H. A. (1972). Theories of bounded rationality. *Decision and organization*, 1(1):161–176.
- Singer, A. C. and Frank, L. M. (2009). Rewarded outcomes enhance reactivation of experience in the hippocampus. *Neuron*, 64(6):910–921.
- Singer, T., Seymour, B., O’Doherty, J., Kaube, H., Dolan, R. J., and Frith, C. D. (2004). Empathy for pain involves the affective but not sensory components of pain. *Science*, 303(5661):1157–1162.
- Singh, P. (2012). Examining the society of mind. *Computing and Informatics*, 22(6):521–543.
- Singh, S., Lewis, R. L., and Barto, A. G. (2009). Where do rewards come from. In *Proceedings of the annual*

- conference of the cognitive science society*, pages 2601–2606. Cognitive Science Society.
- Singh, S., Lewis, R. L., Barto, A. G., and Sorg, J. (2010). Intrinsically motivated reinforcement learning: An evolutionary perspective. *IEEE Transactions on Autonomous Mental Development*, 2(2):70–82.
- Skinner, E. A. (1996). A guide to constructs of control. *Journal of personality and social psychology*, 71(3):549.
- Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological bulletin*, 119(1):3.
- Smallwood, J., Fitzgerald, A., Miles, L. K., and Phillips, L. H. (2009). Shifting moods, wandering minds: negative moods lead the mind to wander. *Emotion*, 9(2):271.
- Smallwood, J., McSpadden, M., and Schooler, J. W. (2007). The lights are on but no one's home: Meta-awareness and the decoupling of attention when the mind wanders. *Psychonomic bulletin & review*, 14(3):527–533.
- Smith, J. (2017). Self-consciousness. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University.
- Smith, J. M. and Price, G. R. (1973). The logic of animal conflict. *Nature*, 246(5427):15.
- Spence, S. A., Brooks, D. J., Hirsch, S. R., Liddle, P. F., Meehan, J., and Grasby, P. M. (1997). A pet study of voluntary movement in schizophrenic patients experiencing passivity phenomena (delusions of alien control). *Brain: a journal of neurology*, 120(11):1997–2011.
- Spira, R. (2017). *The nature of consciousness*. Oxford: Sahaja Publications.
- Spolidoro, M., Baroncelli, L., Putignano, E., Maya-Vetencourt, J. E., Viegi, A., and Maffei, L. (2011). Food restriction enhances visual cortex plasticity in adulthood. *Nature Communications*, 2:320.
- Sridharan, D., Levitin, D. J., and Menon, V. (2008). A critical role for the right fronto-insular cortex in switching between central-executive and default-mode networks. *Proceedings of the National Academy of Sciences*, 105(34):12569–12574.
- Stanovich, K. E. (2004). *The robot's rebellion: Finding meaning in the age of Darwin*. University of Chicago Press.
- Steels, L. (2008). The symbol grounding problem has been solved. So what's next. *Symbols and embodiment: Debates on meaning and cognition*, pages 223–244.
- Stein, D., Kogan, C., Atmaca, M., Fineberg, N., Fontenelle, L., Grant, J., Matsunaga, H., Reddy, Y., Simpson, H., Thomsen, P., et al. (2016). The classification of obsessive-compulsive and related disorders in the ICD-11. *Journal of Affective Disorders*, 190:663–674.
- Sterzer, P., Kleinschmidt, A., and Rees, G. (2009). The neural bases of multistable perception. *Trends in cognitive sciences*, 13(7):310–318.
- Strawson, G. and Watson, G. (1998). Free will. In *Routledge Encyclopedia of Philosophy*. Taylor and Francis.
- Striker, G. (2004). Historical reflections on classical Pyrrhonism and neo-Pyrrhonism. In Sinnott-Armstrong, W., editor, *Pyrrhonian Skepticism*, pages 13–24. Oxford University Press.
- Su, G., Wei, D., Varshney, K. R., and Malioutov, D. M. (2015). Interpretable two-level boolean rule learning for classification. *arXiv preprint, arXiv:1511.07361*.
- Such, F. P., Madhavan, V., Conti, E., Lehman, J., Stanley, K. O., and Clune, J. (2017). Deep neuroevolution: genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning. *arXiv preprint, arXiv:1712.06567*.
- Sun, R., Slusarz, P., and Terry, C. (2005). The interaction of the explicit and the implicit in skill learning: A dual-process approach. *Psychological review*, 112(1):159.
- Sutton, R. S. (1991). Dyna, an integrated architecture for learning, planning, and reacting. *ACM SIGART Bul-*

- letin*, 2(4):160–163.
- Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press: Cambridge, 2nd edition.
- Sutton, R. S., Precup, D., and Singh, S. (1999). Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1-2):181–211.
- Svenaeus, F. (2014). The phenomenology of suffering in medicine and bioethics. *Theoretical medicine and bioethics*, 35(6):407–420.
- Tabor, A., Thacker, M. A., Moseley, G. L., and Körding, K. P. (2017). Pain: a statistical account. *PLoS computational biology*, 13(1):e1005142.
- Taddeo, M. and Floridi, L. (2005). Solving the symbol grounding problem: a critical review of fifteen years of research. *Journal of Experimental & Theoretical Artificial Intelligence*, 17(4):419–445.
- Tagawa, S. (2009). *Living Yogācāra: An introduction to consciousness-only Buddhism*. Simon and Schuster. Translated by Charles Muller.
- Takahashi, Y. K., Batchelor, H. M., Liu, B., Khanna, A., Morales, M., and Schoenbaum, G. (2017). Dopamine neurons respond to errors in the prediction of sensory features of expected rewards. *Neuron*, 95(6):1395–1405.
- Tamar, A., Wu, Y., Thomas, G., Levine, S., and Abbeel, P. (2016). Value iteration networks. In *Advances in Neural Information Processing Systems*, pages 2154–2162.
- Tan, S., Caruana, R., Hooker, G., and Lou, Y. (2017). Detecting bias in black-box models using transparent model distillation. *arXiv preprint, arXiv:1710.06169*.
- Tanaka, K. (1996). Inferotemporal cortex and object vision. *Annual review of neuroscience*, 19(1):109–139.
- Tate, T. and Pearlman, R. (2019). What we mean when we talk about suffering—and why Eric Cassell should not have the last word. *Perspectives in biology and medicine*, 62(1):95–110.
- Teasdale, J. D. (1999). Metacognition, mindfulness and the modification of mood disorders. *Clinical Psychology & Psychotherapy*, 6(2):146–155.
- Teasdale, J. D. and Chaskalson, M. (2011a). How does mindfulness transform suffering? I: The nature and origins of dukkha. *Contemporary Buddhism*, 12(01):89–102.
- Teasdale, J. D. and Chaskalson, M. (2011b). How does mindfulness transform suffering? II: the transformation of dukkha. *Contemporary Buddhism*, 12(1):103–124.
- Teasdale, J. D., Segal, Z. V., Williams, J. M. G., Ridgeway, V. A., Soulsby, J. M., and Lau, M. A. (2000). Prevention of relapse/recurrence in major depression by mindfulness-based cognitive therapy. *Journal of consulting and clinical psychology*, 68(4):615.
- Tejaniya, A. (2008). *Awareness alone is not enough*. Selangor: Auspicious Affinity.
- Teper, R. and Inzlicht, M. (2013). Meditation, mindfulness and executive control: the importance of emotional acceptance and brain-based performance monitoring. *Social cognitive and affective neuroscience*, 8(1):85–92.
- Tesauro, G. (1995). Temporal difference learning and TD-gammon. *Communications of the ACM*, 38(3):58–68.
- Theis, T. N. and Wong, H.-S. P. (2017). The end of Moore’s law: A new beginning for information technology. *Computing in Science & Engineering*, 19(2):41–50.
- Thierry, B., Steru, L., Chermat, R., and Simon, P. (1984). Searching-waiting strategy: A candidate for an evolutionary model of depression? *Behavioral and neural biology*, 41(2):180–189.

- Thrun, S. and Pratt, L., editors (2012). *Learning to learn*. Springer: New York.
- Toda, K. and Platt, M. L. (2015). Animal cognition: monkeys pass the mirror test. *Current Biology*, 25(2):R64–R66.
- Toivonen, H. and Gross, O. (2015). Data mining and machine learning in computational creativity. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 5(6):265–275.
- Tononi, G. and Edelman, G. M. (1998). Consciousness and complexity. *Science*, 282(5395):1846–1851.
- Tossani, E. (2013). The concept of mental pain. *Psychotherapy and psychosomatics*, 82(2):67–73.
- Tsakiris, M., Hesse, M. D., Boy, C., Haggard, P., and Fink, G. R. (2007). Neural signatures of body ownership: a sensory network for bodily self-consciousness. *Cerebral cortex*, 17(10):2235–2244.
- Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131.
- Vago, D. R. and David, S. A. (2012). Self-awareness, self-regulation, and self-transcendence (S-ART): A framework for understanding the neurobiological mechanisms of mindfulness. *Frontiers in human neuroscience*, 6:296.
- Van Boven, L. (2005). Experientialism, materialism, and the pursuit of happiness. *Review of general psychology*, 9(2):132–142.
- Van Gulick, R. (2021). Consciousness. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Winter 2021 edition.
- Van Hateren, J. H. and van der Schaaf, A. (1998). Independent component filters of natural images compared with simple cells in primary visual cortex. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 265(1394):359–366.
- Van Hooft, S. (1998). Suffering and the goals of medicine. *Medicine, Health Care and Philosophy*, 1(2):125–131.
- van Vugt, M. K., Taatgen, N. A., Sackur, J., Bastian, M., Borst, J., and Mehlhorn, K. (2015). Modeling mind-wandering: a tool to better understand distraction. In *Proceedings of the 13th International Conference on Cognitive Modeling*, pages 252–257. University of Groningen Groningen, Netherlands.
- Van Vugt, M. K., van der Velde, M., and ESM-MERGE Investigators (2018). How does rumination impact cognition? A first mechanistic model. *Topics in Cognitive Science*, 10(1):175–191.
- Veehof, M. M., Trompetter, H., Bohlmeijer, E. T., and Schreurs, K. M. G. (2016). Acceptance-and mindfulness-based interventions for the treatment of chronic pain: a meta-analytic review. *Cognitive behaviour therapy*, 45(1):5–31.
- Verduyn, P., Lee, D. S., Park, J., Shablack, H., Orvell, A., Bayer, J., Ybarra, O., Jonides, J., and Kross, E. (2015). Passive facebook usage undermines affective well-being: Experimental and longitudinal evidence. *Journal of Experimental Psychology: General*, 144(2):480.
- Verhaeghen, P. (2017). The self-effacing Buddhist: No(t)-self in early Buddhism and contemplative neuroscience. *Contemporary Buddhism*, 18(1):21–36.
- Vetencourt, J. F. M., Sale, A., Viegi, A., Baroncelli, L., De Pasquale, R., O’Leary, O. F., Castrén, E., and Maffei, L. (2008). The antidepressant fluoxetine restores plasticity in the adult visual cortex. *Science*, 320(5874):385–388.
- Vincent, P., Larochelle, H., Bengio, Y., and Manzagol, P.-A. (2008). Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th International Conference on Machine learning*, pages 1096–1103.

- Vohs, K. D. and Schooler, J. W. (2008). The value of believing in free will: Encouraging a belief in determinism increases cheating. *Psychological science*, 19(1):49–54.
- Von Neumann, J. and Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton University Press.
- Wager, T. D., Atlas, L. Y., Botvinick, M. M., Chang, L. J., Coghill, R. C., Davis, K. D., Iannetti, G. D., Poldrack, R. A., Shackman, A. J., and Yarkoni, T. (2016). Pain in the acc? *Proceedings of the National Academy of Sciences*, 113(18):E2474–E2475.
- Wagner, K., Reggia, J. A., Uriagereka, J., and Wilkinson, G. S. (2003). Progress in the simulation of emergent communication and language. *Adaptive Behavior*, 11(1):37–69.
- Walter, S. (2009). Epiphenomenalism. In *Encyclopedia of neuroscience*, pages 1137–1139. Springer.
- Watkins, E. and Teasdale, J. D. (2004). Adaptive and maladaptive self-focus in depression. *Journal of affective disorders*, 82(1):1–8.
- Wegner, D. M. (2002). *The illusion of conscious will*. Cambridge, MA: MIT Press.
- Wegner, D. M. (2003). The mind's best trick: how we experience conscious will. *Trends in cognitive sciences*, 7(2):65–69.
- Weike, A. I., Schupp, H. T., and Hamm, A. O. (2007). Fear acquisition requires awareness in trace but not delay conditioning. *Psychophysiology*, 44(1):170–180.
- Weiss, K., Khoshgoftaar, T. M., and Wang, D. (2016). A survey of transfer learning. *Journal of Big data*, 3(1):9.
- Wen, W. and Haggard, P. (2018). Control changes the way we look at the world. *Journal of cognitive neuroscience*.
- Westbrook, R. F., Iordanova, M., McNally, G., Richardson, R., and Harris, J. A. (2002). Reinstatement of fear to an extinguished conditioned stimulus: two roles for context. *Journal of Experimental Psychology: Animal Behavior Processes*, 28(1):97.
- Whiten, A., McGuigan, N., Marshall-Pescini, S., and Hopper, L. M. (2009). Emulation, imitation, over-imitation and the scope of culture for child and chimpanzee. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1528):2417–2428.
- Whitmer, A. J. and Gotlib, I. H. (2013). An attentional scope model of rumination. *Psychological bulletin*, 139(5):1036.
- Wiech, K., Ploner, M., and Tracey, I. (2008). Neurocognitive aspects of pain perception. *Trends in cognitive sciences*, 12(8):306–313.
- Wielgosz, J., Goldberg, S. B., Kral, T. R., Dunne, J. D., and Davidson, R. J. (2019). Mindfulness meditation and psychopathology. *Annual review of clinical psychology*, 15:285–316.
- Wikenheiser, A. M. and Redish, A. D. (2015). Hippocampal theta sequences reflect current goals. *Nature Neuroscience*, 18(2):289–294.
- Williams, J. M. G. (2008a). Mindfulness, depression and modes of mind. *Cognitive Therapy and Research*, 32(6):721.
- Williams, P. (2008b). *Mahayana Buddhism: The Doctrinal Foundations*. Routledge.
- Wolpert, D. M., Miall, R. C., and Kawato, M. (1998). Internal models in the cerebellum. *Trends in cognitive sciences*, 2(9):338–347.
- Wong, R. O. (1999). Retinal waves and visual system development. *Annual review of neuroscience*, 22(1):29–47.
- Wright, A. (2011). A criticism of the IASP's definition of pain. *Journal of Consciousness Studies*, 18(9-10):19–44.
- Wright, R. (1994). *The Moral Animal*. Vintage Books.

- Wright, R. (2017). *Why Buddhism is true: The science and philosophy of meditation and enlightenment*. Simon and Schuster.
- Xu, D., Li, T., Li, Y., Su, X., Tarkoma, S., Jiang, T., Crowcroft, J., and Hui, P. (2020). Edge intelligence: Architectures, challenges, and applications. *arXiv preprint, arXiv:2003.12172*.
- Yehuda, R. (2002). Post-traumatic stress disorder. *New England journal of medicine*, 346(2):108–114.
- Yi, K., Wu, J., Gan, C., Torralba, A., Kohli, P., and Tenenbaum, J. (2018). Neural-symbolic VQA: Disentangling reasoning from vision and language understanding. In *Advances in Neural Information Processing Systems*, pages 1031–1042.
- Yoshida, W., Seymour, B., Koltzenburg, M., and Dolan, R. J. (2013). Uncertainty increases pain: evidence for a novel mechanism of pain modulation involving the periaqueductal gray. *Journal of Neuroscience*, 33(13):5638–5646.
- Zeng, J., Ustun, B., and Rudin, C. (2017). Interpretable classification models for recidivism prediction. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 180(3):689–722.
- Zenke, F., Gerstner, W., and Ganguli, S. (2017). The temporal paradox of hebbian learning and homeostatic plasticity. *Current opinion in neurobiology*, 43:166–176.
- Zhang, K. and Sejnowski, T. J. (2000). A universal scaling law between gray matter and white matter of cerebral cortex. *Proceedings of the National Academy of Sciences*, 97(10):5621–5626.
- Zhang, R., Brennan, T. J., and Lo, A. W. (2014). The origin of risk aversion. *Proceedings of the National Academy of Sciences*, 111(50):17777–17782.
- Zhou, B., Zhao, H., Puig, X., Xiao, T., Fidler, S., Barriuso, A., and Torralba, A. (2019). Semantic understanding of scenes through the ADE20K dataset. *International Journal of Computer Vision*, 127(3):302–321.
- Zhuang, C., Yan, S., Nayebi, A., Schrimpf, M., Frank, M. C., DiCarlo, J. J., and Yamins, D. L. (2021). Unsupervised neural network models of the ventral visual stream. *Proceedings of the National Academy of Sciences*, 118(3).
- Ziemke, T. (2007). The embodied self: Theories, hunches and robot models. *Journal of Consciousness Studies*, 14(7):167–179.
- Zinkevich, M., Johanson, M., Bowling, M., and Piccione, C. (2008). Regret minimization in games with incomplete information. In *Advances in neural information processing systems*, pages 1729–1736.
- Zinkevich, M., Weimer, M., Li, L., and Smola, A. J. (2010). Parallelized stochastic gradient descent. In *Advances in neural information processing systems*, pages 2595–2603.